Escaping Flatland: Designing data visualizations in augmented reality

Peter Andringa

Senior Honors Thesis
Hussman School of Journalism and Media
University of North Carolina at Chapel Hill

Chapel Hill
April 7, 2020

Approved:

____

Professor Steven King, Advisor

____

Professor Laura Ruel, Reader

© 2020
Peter Andringa
ALL RIGHTS RESERVED

Abstract

Peter Andringa: Escaping Flatland: Designing data visualizations in augmented reality
(Under the direction of Steven King)

Augmented reality is becoming increasingly common in digital news, as nearly every major newsroom has released AR-driven stories in recent years. For this technology to become more than a fad, AR creators will need to pay attention to the user experience and find ways of telling stories that people will understand and remember. This study aims to solve a small part of that challenge, through a user-testing study of data visualizations in augmented reality. Drawing upon the 2D data visualization theories of Edward Tufte, Jaques Bertin, and others, three visualizations were created in AR and examined using user testing methodologies set out by previous AR researchers, as well as user experience researchers like Nielsen and Norman. Twenty-three undergraduate students viewed the augmented reality in the context of a sample news article, then answered comprehension questions and were interviewed about their experiences. Their reactions showed that AR design has particular complications to be careful of: users frequently expressed confusion about objects being occluded behind each other, labels being hard to read, and scales being difficult to measure. All their observations were condensed into five key best practices for AR development: “Keep all items in view,” “Optimize for macro/micro compositions,” “Assume users won’t move around,” “Use labels sparingly,” and “Allow opportunities to learn and fail.” With these guidelines, future development of AR data visualizations can continue to improve user experiences and help promote more widespread adoption of this new technology.

Acknowledgements

Even a medium-scale endeavor like this one owes thanks to many people who helped it along the way. I owe huge gratitude to my advisor Steven King and reader Laura Ruel, who have guided my exploration of news technology over the years and opened doors to countless opportunities I would have never imagined, including this thesis itself. Similarly, I want to thank Barbara Friedman and many other faculty in the Hussman School (and Duke's DeWitt Wallace Center), who gave me a love for journalism and guided my writing over the years.

I have learned and grown thanks to the staff and leadership of the Robertson Scholars Program, which has shaped my college career (and beyond) in more ways than I can count. I should also thank the many friends who accompanied me on hours-long stints in Chapel Hill coffeeshops this year, offering motivation and much-needed breaks while writing.

Finally, I want to acknowledge and thank the devotion of my parents and family, without whom I would not be here today, and whose love and encouragement guide me all the time.

Table of Contents

Introduction

In 1990, Edward Tufte recognized the difficulty of designing in 2D. "Even though we navigate daily through a perceptual world of three spatial dimensions, the world portrayed on our information displays is caught up in the two-dimensionality of the endless flatlands of paper and video screen," the statistician wrote in one of his influential books on data visualization, *Envisioning Information*. As he puts it, "Escaping this flatland is the essential task of envisioning information."[1] Now considered one of the pioneers of modern data visualization, Tufte's work has provided a philosophy and set of best practices for displaying complex data in clear illustrations and graphics.

More than thirty years after their publication, the fundamentals of Tufte's theories are applied in multiple fields, but perhaps no use is as widely-consumed as the news media. Modern newsrooms employ dozens of visual journalists charged with translating complex news into digestible graphics, often using the same charting best practices espoused by Tufte. These graphics are aired on television, printed in newspapers, and hosted on news websites originating in nearly every country in the world.

The growth of news graphics as a field is partially tied to the growth of online news, because the internet offers new methods of interaction and visual journalism never before possible in print or television. Computing also opens up the range of data that can be visualized, allowing anyone with a laptop computer and Microsoft Excel to construct charts that illustrate millions of data points. The largest newsrooms have even developed "interactive news" desks: teams of journalists equipped with software-engineering skills who can build "news apps"

[1] Edward Tufte, *Envisioning Information*, 12.

that visualize complex databases or allow for extensive user interaction.[2] The stories these teams generate are frequently the most-read across the entire newsroom, generating engagement far beyond what a text-only story might usually engender.[3]

Yet, data visualization on the web is still fundamentally the same as when Tufte first wrote about in 1990; restricted to the "flatland" of 2D space offered by page and screen alike. Data visualization is still constrained by these lack of dimensions: as Tufte notes, "nearly every escape from flatland demands extensive compromise, trading off one virtue against another."[4]

New technologies for augmented reality (AR) offer potential to break these constraints and move beyond flatland. AR — which lets users view digitally-constructed objects in real, 3D space around them — has previously been restricted to the confines of sci-fi holograms and fictional spy glasses with built-in displays. However, in recent years augmented reality has come significantly closer to widespread use, through experimental headsets with "goggles" that users can put on and, crucially, the development of smartphone-based AR applications. Now, nearly every smartphone user has the capability to view augmented reality: by pointing their phone at the room around them, virtual objects can be superimposed into it through their phone screen. While this experience is still far from the glasses-based augmented reality that futurists believe will eventually be part of our daily lives,[5] it still provides a taste of a future that seamlessly blends the digital world with the 3D world around us.

---

[2] Nikki Usher, *Interactive Journalism*, 18.

[3] Nikki Usher, *Interactive Journalism,* 63.

[4] Tufte, *Envisioning Information*, 15.

[5] Peddle, *Augmented reality: where we will all live*, 5.

Newsrooms have been some of the earliest adopters of AR technology, eager to escape the bounds of 2D flatland and experiment with storytelling in 3D space. Prominent newspapers like the New York Times and The Washington Post have incorporated it into their apps which are already installed on millions of phones, allowing readers to explore the Apollo 11 landing site or a Winter Olympics speed-skating event in 3D. Other media organizations have released their own dedicated AR apps, or built partnerships that allow both their journalists and their readers to try out the new format.

For now, most of these experiments use a technique called photogrammetry to capture real-world objects using hundreds of photos at different angles, composing them all into a 3D object that can be viewed in AR. This is sort of a 3D analogue of photography or videography, providing a true-to-life representation of the real world in a new medium. Data visualization (using abstract images or shapes that highlight trends in data, instead of real-life photos) is less common in news AR projects. This may be because it is more time and labor-intensive compared to photogrammetric models, or because users are so unaccustomed to it that the technology can sometimes distract from the underlying story.

Newsrooms also might shy away from data visualization in augmented reality because current interfaces for AR are somewhat clunky and hard to use. Viewers have to peer through the small window of their phone screen, constantly holding it up and moving around to see an entire shape or graphic. This limits the amount of time users want to spend exploring a data visualization in AR, and also adds one more layer of complexity to what could be an already-complex 3D data visualization.

Should high-fidelity, glasses-based AR become more widespread, they would be an obvious way for graphics designers and data journalists to represent complex datasets more clearly without the constraints of smartphone AR. Moving from 2D to 3D data visualizations unlocks a whole new dimension for graphics to take advantage of, resolving some of the problems Tufte identified about flatland. At that point, newsrooms will need a clear set of best practices and design patterns that work for presenting data in AR — a 3D update to the already widespread best practices of 2D data visualization.

This thesis explores possible uses of data visualization in augmented reality and identifies how those established best practices may or may not apply to 3D space. Through a series of AR experiments and user testing, the findings aim to further the development of new types of data visualization for storytelling, and prepare for a world where 3D interfaces could become the primary way media consumers encounter the news.

Background: Origins of Data Visualization

Because they seem instinctual to modern readers, it's easy to assume that the bar chart — and the larger idea of visualizing numbers with geometry — is an ancient creation, maybe emerging in the time of Pythagoras and the Greek mathematicians. But in fact, even the most basic forms of data visualization are a modern phenomenon, created in late-eighteenth century United Kingdom to help the global power manage its growing economy and colonial operations.[6]

Previously, the closest analogue to data visualization was cartography, which requires the mapmaker to decide how to express elements of the real world in an imaginary 2-dimensional representation. Early mapmaking was equally art and science, with cartographers and their

---

[6] Tufte, *Visual Display of Quantitative Information*, 32-34.

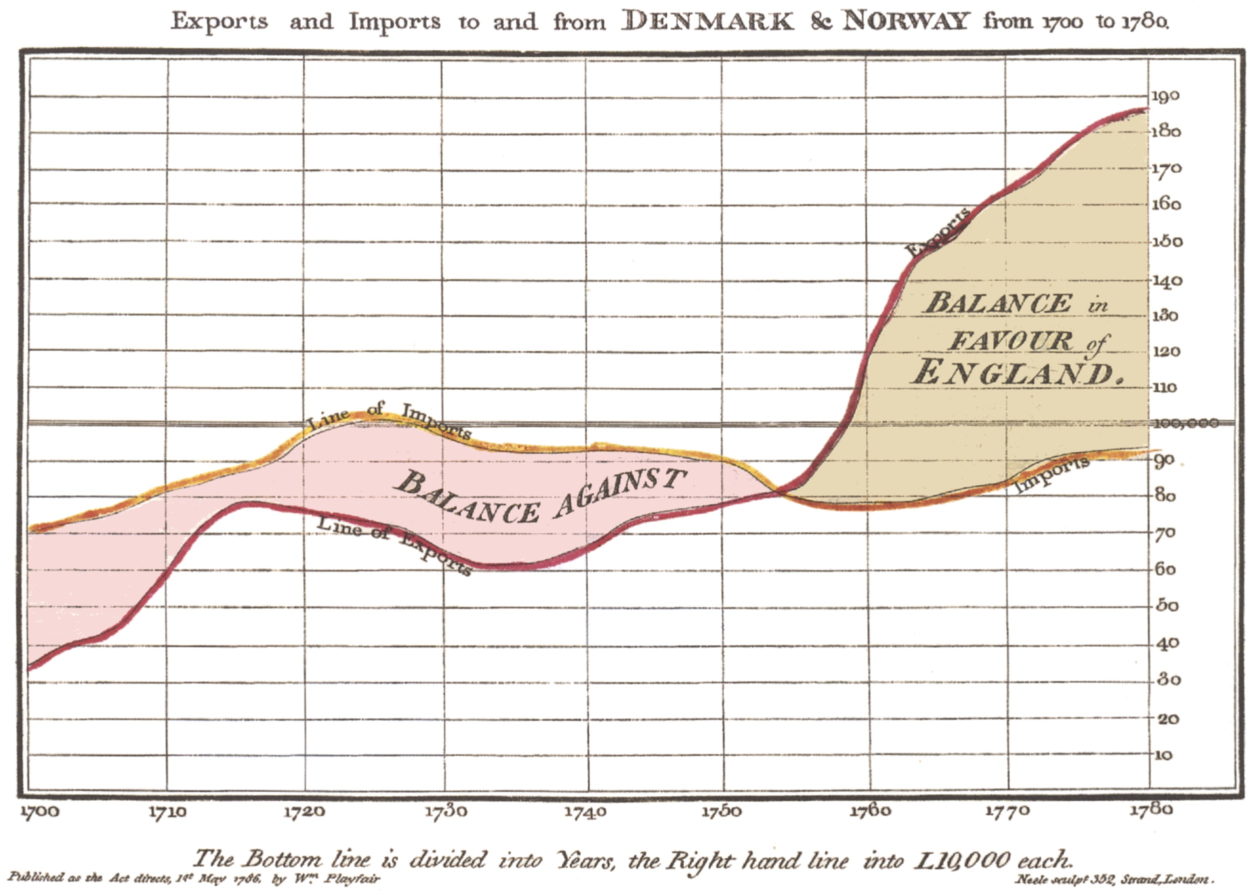

Figure 1: A time series chart showing England's foreign trade. Drawn by William Playfair, 1786 in the *Commercial and Political Atlas*

nations racing to create the most beautiful and precise maps to aid their navigation. Still, maps are ultimately just a representation of the real world — the next revolution would come with the idea of representing abstract *ideas* (that is, numbers) on paper in the physical world. The credit for that innovation usually goes to William Playfair, a Scottish political economist who advised the British government on its foreign relations. Playfair invented the line chart, bar graph, and area chart in his *Commercial and Political Atlas* published in 1786 (fig. 1), and later introduced a form of pie chart in *Statistical Breviary,* published in 1801.[7]

[7] Tufte, *Visual Display of Quantitative Information*, 32-34.

The industrial revolution and its processes generated vast quantities of data and created the modern field of statistics, which continued the need for new forms of graphical representations for data into the nineteenth century. Florence Nightingale, best known as the founder of modern nursing, was another influential pioneer who invented the grouped bar chart and the radial area chart for use in publications that sought to explain Crimean War-era British Army deaths due to disease (fig. 2).[8] American engineer Willard Brinton later developed other well-known chart types, like slope graphs and small multiples, in his 1914 business book, *Graphic Methods for Presenting Facts.*

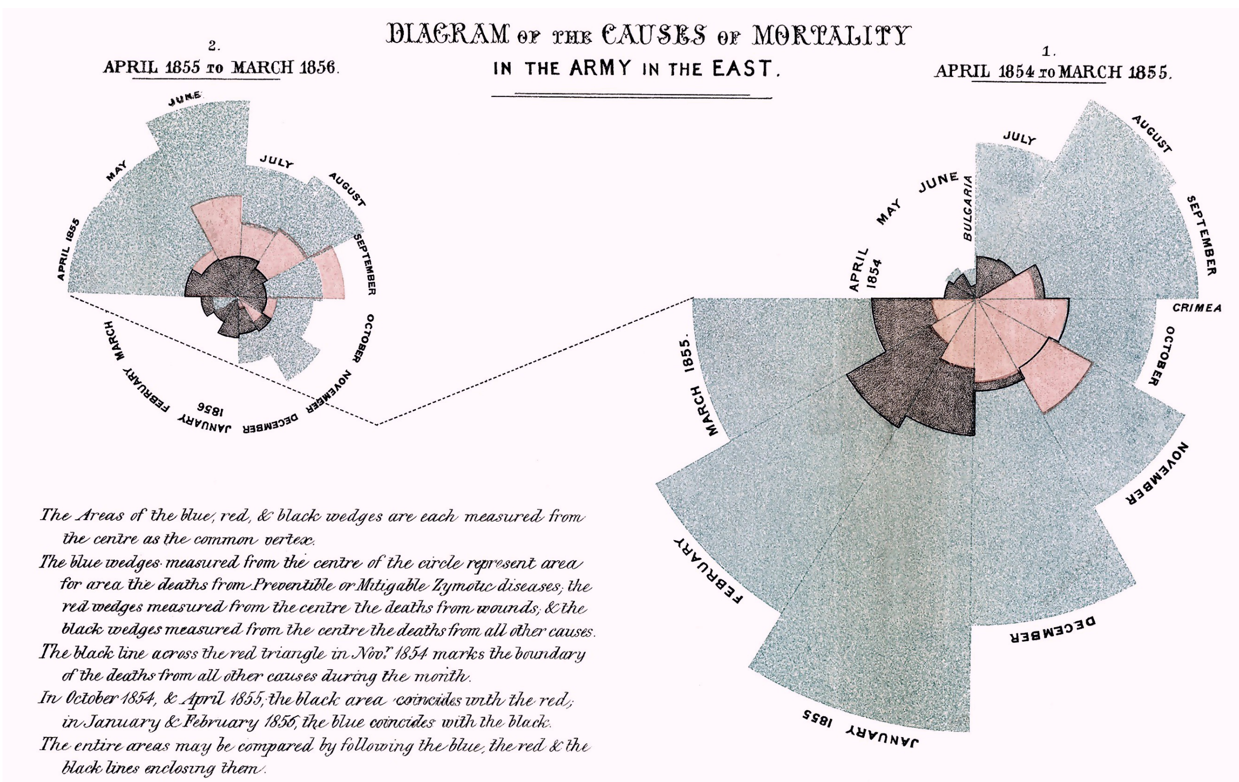

Figure 2: A radial chart showing causes of death for British soldiers during the Crimean War. Drawn by Florence Nightingale, 1858.

Perhaps the most influential theorist in data visualization today is Edward Tufte, a Yale statistician and designer whose books are *de rigueur* for business analysts, journalists, and other

[8] Berinato, *Good Charts*, 18.

data-graphics professionals. His theories, presented in 1983's *The Visual Display of Quantitative Information* and 1990's *Envisioning Information,* set out new chart types, provided guidelines for designing effective visuals, and proposed thoughts on the future of data visualization. Tufte's theories continue to shape the design of graphics, even as computers and the internet have made it substantially easier to create them. Where Tufte formerly had to measure angles or draw lines by hand, digital graphics tools can draw millions of lines in fractions of a second, enabling more complex and more interactive forms of data visualization than ever before.

One of the most recently developed forms of computer graphics is augmented reality, which projects digital content into real-life space. The concept of augmented reality has figured in popular science fiction for decades, including Phillip K. Dick's book *Minority Report* (1956) and Alan Dean Foster's movie *Star Wars* (1977), as two examples.[9] While it existed as a compelling science-fiction plot, the term "augmented reality" wasn't coined until 1992, in a paper by Boeing researchers Thomas Caudell and David Mizell detailing their creation of a head-mounted display to provide instructions for manufacturing workers.[10] In the same decade, other researchers created head-mounted displays for virtual reality (VR), which places users in an immersive 3D space instead of inserting digital content into the 3D world around them.

Today, immersive technology is widespread and represented in many formats. VR headsets, like Facebook's Oculus brand,[11] are widely available and often marketed for gaming. AR glasses are in active development (e.g., Microsoft's Hololens, targeted at the business

---

[9] Peddle, *Augmented reality: where we will all live*, 1.

[10] Caudell and Mizell, "Augmented Reality: An Application of Heads-Up Display Technology to Manual Manufacturing Processes."

[11] "Oculus, from Facebook" (https://www.oculus.com)

market,[12] or the Magic Leap, which is slightly more consumer-focused[13]), but many of them are bulky, slow, and offer a limited field-of-view that makes them less effective.

However, a different form of augmented reality is already in the pockets of millions of consumers: iOS and Android have sophisticated platforms that allow smartphones to double as AR viewers, overlaying digital objects into the real world using the camera as a viewfinder. A smartphone certainly does not offer an ideal experience: peering through a 6-inch screen is not very immersive, and one's hands tire and cramp holding it up for a long time. Still, smartphones offer AR content to a significantly wider audience, and also allow creators to experiment with new designs — promoting the development of tools and techniques that will be transferrable to more advanced AR technologies in the future.

---

[12] "HoloLens 2: Mixed Reality Technology for Success" (https://www.microsoft.com/en-us/hololens)

[13] "Magic Leap" (https://www.magicleap.com)

Literature Review

Design theories of data visualization

Since data visualization is equal parts art and science, it's useful to look to theories and best practices established by experienced designers and academics who study information design. In *Envisioning Information*, Tufte defends his own practice with a similar argument, writing, "Ideas not only guide work, but also help defend our designs (by providing reasons for choices) against arbitrary taste preferences."[14] In his book, Tufte sets out a number of principles for his "Theory of Data Graphics," focused primarily on an ethos of simplifying and making graphics as information-dense as possible. Many of his maxims will be useful for this study, including:

- "Maximize the data-ink ratio" by eliminating unnecessary elements and ensuring that each element of the graphic represents the data in some way.[15]
- "Forgo chartjunk," defined as illustrations or ornamental decoration that don't relate to the data itself. Tufte attacks common examples, like 3-D perspective charts, illustrative art that incorporates the chart, or special hatching or shading (fig. 3).[16]
- "Mobilize every graphical element… to show the data" by using lots of possible dimensions to encode data, making every variable tied to a piece of information and allowing readers to gain more from a closer reading.[17]

---

[14] Tufte, *Envisioning Information*, 82.

[15] Tufte, *Envisioning Information,* 105.

[16] Tufte, *Envisioning Information,* 121.

[17] Tufte, *Envisioning Information,* 139.

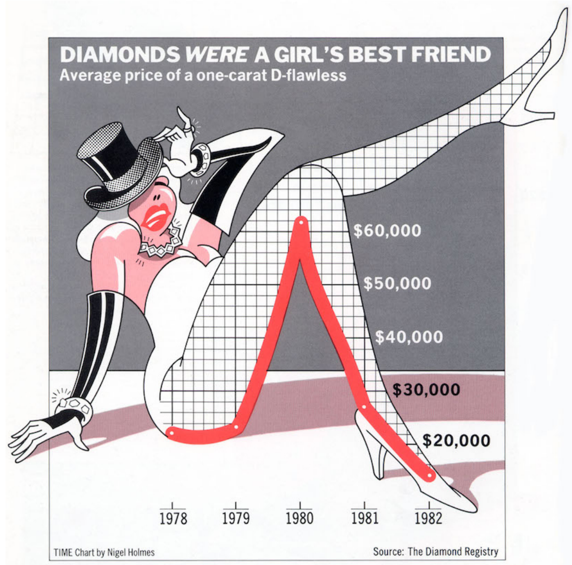

Figure 3: An example of "chartjunk" that Tufte cites in *Envisioning Information*. ("A Gem That Lost Its Luster," Time 120, no. 9, August 30, 1982)

Many of Tufte's suggestions build on the earlier visualization theory of Jacques Bertin, a French cartographer who published *Semiologie Graphique* in 1967. Bertin was one of the first to articulate the concept of "visual variables" that can encode information in graphics — position, size, shape, color, brightness, orientation, and texture — emphasizing "expressiveness" and "effectiveness" overall.[18] Ultimately however, theories and recommendations regarding data visualization are valuable only if they contribute to producing clear, informational, and even beautiful graphics.

Quantitative task-based visualization evaluations

Most empirical studies of data visualizations have sought to understand whether a data graphic is truly effective for its audience. A number of studies use a "tasks-to-be-done" framework to examine data graphics, testing whether readers viewing a given graphic are able to

---

[18] Berinato, Scott, *Good Charts,* 21.

quickly and accurately solve a task like determining the distance between two points or the relative size of two variables.

This technique was largely pioneered by William Cleveland and Robert McGill's 1984 study to understand how readers interpreted visualizations, a theory that they called "graphical perception."[19] In that study, the authors divided their graphics into "elementary tasks" that are testable through experimentation: things like position, length, direction, angle, area, volume, curvature, and shading, similar to the "visual variables" articulated by Bertin. By measuring the accuracy of each of these variables, they reasoned, they might begin to create empirical standards of data visualization design. The authors asked subjects to determine the largest or smallest point, compare two lengths, or evaluate the share of a pie chart. Ultimately, the study found that humans are better able to quickly compare position — say, between two dots positioned left or right — than length of two lines, or angles of two pie chart slices. That allowed the authors to suggest a number of revisions to common chart types, like replacing pie charts with bar charts when possible.

In 2010, Jeffrey Heer and Michael Bostock replicated Cleveland and McGill's studies of graphical perception, with updates for the internet age: digital graphics and crowd-sourced participants from Amazon's Mechanical Turk.[20] Their findings mirrored those of Cleveland and McGill, suggesting the new internet-based forms of study are an effective tool for twenty-first century researchers. Heer and Bostock went on to test new forms of graphics using the same digital methodology. Similar to prior studies, they found that viewers were better able to judge

[19] Cleveland and McGill, "Graphical Perception"

[20] Heer and Bostock, "Crowdsourcing Graphical Perception."

rectangles' areas than circles. They also found that the screen quality affects the perception of opacity, grid spacing, and color for digital graphics — an important finding that hardware constraints may affect perception, now that many graphics are distributed to varied devices via the internet.

The fallibility of human perception has at times posed problems for the data visualization designer, so other studies have explored ways to overcome it. Some designers have suggested a possible solution by encoding data in a redundant way; that is, by using more than one "visual variable" to express the same datapoint.[21] Tufte criticized this practice in *Visual Information* as a waste of space that made visualizations too complex. But what if it also aided in comprehension? Russel Chun set out to determine this in a 2017 experiment,[22] in which he replicated Cleveland and McGill's study but added an additional variable (color) to represent values in addition to size. Chun found no discernible improvement from the redundant encoding, which suggests that humans pay attention to one visual variable at a time, and that those variables might be better put to use expressing a new visual dimension of the data.

Other studies go further up the abstraction ladder, studying not only the individual variables of visualization, but how readers perceive and become accustomed to more complicated graphics. Some studies have shown that pattern recognition is a significant element of reading graphics: Shah et al.[23] noted that audiences have been trained to understand line graphs up-and-to-the-right as increasing, and down-and-to-the-right as decreasing — even though that's only a convention of placing graph axes.

---

[21] Aisch et. al, "Data Stories Podcast: Episode 7 | Color."

[22] Chun, "Redundant Encoding in Data Visualizations."

[23] Shah et al., "Graphs as aids to knowledge construction."

Some scholars and designers have studied the potential of pattern-recognition by using "visual metaphors" in visualizations: repeatable designs that represent the same relationship to the real world. Ziemkiewicz and Kosara, for example, presented users with either area diagrams (which appear to be a nested set of squares) or node-link trees (where shapes are linked together into a larger tree-like shape, to show relationships) and posed questions that primed them with a verbal metaphor of containers or tree branches, respectively. They found that priming viewers with the right visual metaphor improved the speed and accuracy of readers answers, but the most-important factor was whether users internalized those metaphors over time.[24] Ziemkiewicz and Kosara concluded:

> This would suggest that more immediately successful visualizations are likely to be those that match the user group's existing metaphors about their data and the work they need to perform with it, and that the discovery of these existing metaphors should be an important part of the visualization design process.

Tufte also promotes the same concept in *Visual Display*, recommending what he calls "small multiple" graphic designs[25]: creating a visual metaphor that can be re-used multiple times throughout a visualization, so that readers can spend time internalizing the metaphor on one small example, then see it recurring in multiple later cases.

This offers a synthesis of some of the seminal quantitative studies of data visualization. There are many more of these types of studies — enough for a robust meta-analysis covering more than three dozen studies from just nine years, 1991-2000, which found a wide mix of statistically significant results that often contradicted each other,[26] which was attributed to the

---

[24] Ziemkiewicz and Kosara, "The Shaping of Information by Visual Metaphors," 1275.

[25] Tufte, *Visual Display*, 170.

[26] Chen and Yu, "Empirical studies of information visualization: a meta-analysis," 851.

diverse range of methodologies and lack of standardized tasks to examine. Despite this confusion, the volume of work speaks to the interest and the unknown in this still-young field.

Qualitative and human-centered visualization evaluation

In other research, scholars have taken a qualitative approach to understanding data visualization.[27] Their holistic approach seeks to examine the phenomenon of graphics, rather than the quantitative measurement of their effects.

Meta-analyses of the field have sought to "encourage an approach to evaluation that is based on evaluation goals and questions instead of methods, and to encourage… a wider view on the possibilities for evaluation in data visualization."[28] That 2012 meta-analysis, by led Heidi Lam, coded more than 860 data-visualization studies into scenarios based on the purpose of each study: work environments, visual data analysis, communication through visualization, collaborative data analysis, user performance, user experience, and visual algorithms. The authors found that 85% of the literature focused on the three scenarios most closely tied to the abstract visualization itself (user experience, user performance, and visual algorithms), and 15% on the more comprehensive "process" studies that examine the comprehensive effect of the visualizations. They ask, "Are we as a community less welcoming to these different — often qualitative — types of evaluations?"[29]

Since then, a number of scholars have tried to formulate guidelines for evaluating visualizations in this new manner. Many have drawn the larger discipline of human-computer interaction, particularly its user-focused theory of "Human-Centered Design." While Human-

---

[27] Stasko, "Value-Driven Evaluation of Visualizations," 48.

[28] Lam et al., "Empirical Studies in Information Visualization: Seven Scenarios," 1522.

[29] Lam et al., "Empirical Studies in Information Visualization: Seven Scenarios," 1532.

Centered Design is usually applied to product design, Frietas et al. argue that the concepts are also useful for designing visualization methodologies. They offer the following guidelines for human-centered evaluation of visualizations:[30]

> 1. The context of usage for evaluation must be defined before the beginning of evaluation.
> 2. Evaluation needs to know who the users are, of what are their goals, and to decide which users to support.
> 3. Evaluation needs to understand which tasks users need to perform and their characteristics (steps, constraints, and other tasks attributes like frequency, priority, etc.) and to decide which tasks to support.
> 4. Evaluation should be performed earlier in the design-development cycle.

Following these guidelines, they argue, leads to a more comprehensive measurement of visualization effectiveness than does looking at individual components of visualizations separately.

An important dimension of visualization effectiveness is interpretation: whether a graphic correctly communicates the main themes it was intended to highlight. Shah et al. had participants narrate their interpretation of graphics from social-science textbooks, some of which had been redesigned with the same data in different formats.[31] Those narrations revealed that readers generally failed to understand the point of the original graphs, and that redesigns could improve understanding — but more importantly, that readers' interpretation was closely tied to the format of graph (e.g., bar or line) and their own recognition of similar formats and patterns in the past.

Other studies focus on persuasion: whether charts are an effective tool for convincing the reader to change their mind. A study by Pandey et al., for example, found that graphics were more persuasive than tables, when each was presented to readers alongside a news story about a

---

[30] Frietas et al., "User-Centered Evaluation of Information Visualization Techniques," 329.

[31] Shah et al., "Graphs as aids to knowledge construction."

controversial social or political stance — but only when the reader didn't have a strong preexisting opinion on the issue. In that case, the tables actually had a stronger effect, which the authors hypothesized might be due to the impression that the table communicated "concrete" numbers.[32]

Another attribute of chart communication is memorability: whether charts increase or decrease the likelihood that a fact sticks in readers minds. Borken et al. conducted surveys testing 410 charts from governments, news organizations, and other publishers, which they classified into a taxonomy of different designs. The authors found that more complex charts like scatterplots and tree diagrams were actually more memorable than simple lines or bars, perhaps because they are less common in the real world.[33] They also found that charts with lots of color or some form of illustration were more memorable, contradicting the general advice of Tufte and other proponents of simple charts, free of "chartjunk."

Many studies of illustrated charts have come to the same conclusion: while generally eschewed by Tufte and other statistics professionals, illustrated and colorful charts appear to aid in memory retention[34] and might improve understanding the concepts behind graphics,[35] especially in time-constrained environments.[36] However, those embellishments actually impeded readers when their task was to search for specific information in a graphic, suggesting that highly

---

[32] Pandey et al., "The Persuasive Power of Data Visualization," 2219.

[33] Borken et al., "What Makes a Visualization Memorable?" 2314.

[34] Bateman et al., "Useful junk?"

[35] Tankard, "Effects of Chartoons and Three-Dimensional Graphs"

[36] Borgo et al., "An Empirical Study on Using Visual Embellishments," 2767.

decorated charts might be useful for publishing and mass communication, but not for detailed reference or analysis.[37]

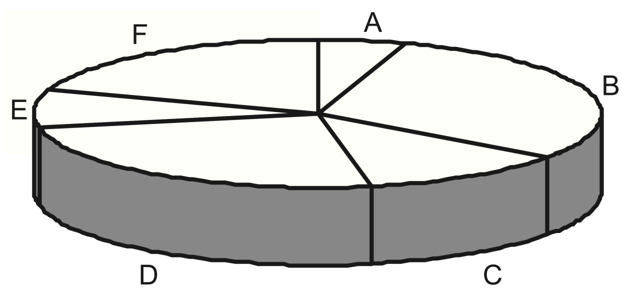

Figure 4: An example of a 3D-perspective pie chart that can confuse readers' perceptions of relative size. (Rangecroft, "As Easy as Pie")

Another variety of Tufte's chartjunk are 3-D perspective graphs, which take a familiar bar or pie chart and try to make it more interesting by giving it a skewed perspective in relation to the reader. Some of the same studies mentioned above[38] as well as other studies of 3D pie charts (fig. 4)[39] have shown that that perspective has the possibility of introducing confusing visual effects and optical illusions, ultimately making it more difficult for the reader to compare various lengths, areas, or shapes. As such, the consensus view of research seems to align with Tufte, agreeing that perspective-drawn 3D charts are best excluded from data graphics.

Data visualization in 3D immersive (AR/VR) environments

But what if 3D graphics were to be drawn in 3D space? The increasingly widespread nature of augmented reality headsets and phone-based AR has dramatically increased the amount

---

[37] Borgo et al., "An Empirical Study on Using Visual Embellishments."

[38] Tankard, "Chartoons and Three-Dimensional Graphs"

[39] Rangecroft, "As Easy as Pie," 426

of research into interfaces and data visualization in augmented reality, as researchers are scrambling to find out what works and what doesn't.

In a 2008 review of the recent and potential applications for 3D visualization, Ifan Shepherd notes that 3D poses a few attractive advantages for displaying data: it allows for an extra dimension, it's more natural and intuitive to the real world, and it allows for multiple layers of map data that don't obscure each other.[40] However, he also notes many issues like the fact that 3D perspective can distort our perception of distance and size, and that object occlusion is still a problem in 3D environments. Ultimately, he suggests that 3D visualizations might *not* be inherently better, which aligns with multiple studies which found no significant difference between 2D and 3D graphics[41] and others that found a negative effect due to perspective distortion.[42] Shepherd concludes that "virtual 3D worlds are sometimes more challenging than the real thing."

However, Shepherd's conclusion is based on the desktop-based 3D technology that existed in 2008, while the intervening decade has brought transformational changes to the user experience of 3D visuals. No longer is 3D content slow and clunky; modern computers and even phones can render complex 3D graphics in less than 1/60th of a second.[43] Additionally, previous studies have shown that one of the biggest impediments to comprehension of desktop-based 3D

---

[40] Shepherd, "Travails in the Third Dimension," 204.

[41] Bleisch and Dykes, "Quantitative data graphics in 3D desktop- based virtual environments," 627; Seipel and Carvalho, "Solving Combined Geospatial Tasks Using 2D and 3D," 6.

[42] Lind et al., "Metric 3D structure in visualizations," 54; Bleisch, "Toward Appropriate Representations of Quantitative Data in Virtual Environments."

[43] "Apple Developer: SceneKit PreferredFramesPerSecond" (https://developer.apple.com/documentation/scenekit/scnview/1621205-preferredframespersecond)

graphics is the difficulty of interaction,[44] and virtual reality and augmented reality allows for a much more natural interaction pattern of movement than older mouse-driven desktop interfaces. As such, different research is required to understand the nuances of data visualization in immersive 3D environments.

Of the few studies in this field, most focus on virtual reality, technology which has been available for slightly longer and has headsets more widely available to consumers. One study by Yang et al. looked at maps in VR, seeking to understand how users might interact with geographic data in a 3D environment, for example. The researchers examined three types of maps: flat, curved, spherical, and egocentric (in which the viewer is *inside* the sphere) and examined task-completion times like locating a specific city or country in each one. They discovered that the spherical globe was the most effective and well-liked by the users (perhaps because it is familiar and takes advantage of the 3D environment), whereas the egocentric globe was also effective but caused some participants to report motion-sickness.[45]

Other studies looked at mobile-device based augmented reality, the most widely usable format of AR today. Many of these studies experiment with node-link graphs in AR, a format sometimes confusing to view in 2D. Some of those studies examined data encoding along the link edges, finding that 3D position and color were effective variables[46] for encoding data, and that visualizations are better when they are simple enough that they avoid visual confusion with

---

[44] Herman et al., "Evaluation of User Performance in Interactive and Static 3D Maps."

[45] Yang et al., "Maps and Globes in Virtual Reality."

[46] Belcher et al., "Using Augmented Reality for Visualizing Complex Graphs in Three Dimensions."; Büschel et al., "Investigating Link Attributes of Graph Visualizations in Mobile Augmented Reality."

the background or occlusion from other objects.[47] However, this group of studies was mostly task-driven, measuring accuracy and completion rates as opposed to more subjective interpretations and comprehension of the visualization. Büschel et al. also examined the movement of users while viewing the graph in virtual reality, and found that users tended to walk around the object to view it from all areas, suggesting that an "outside-in" model of perception is preferred to an environment where the user is placed *within* the visualization itself.[48]

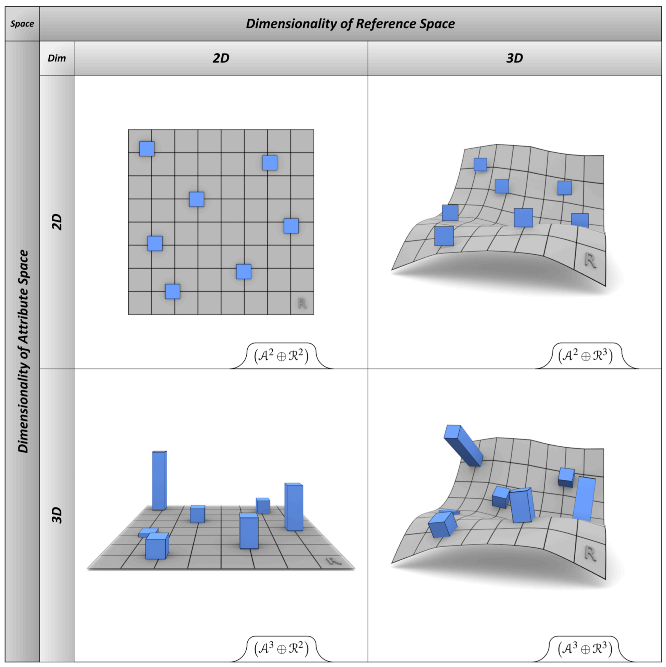

Figure 5: Dübel et al.'s categorization of 3D data graphics based on their dimensionality.

In a novel, systematic categorization of different varieties of 3D visualizations, Dübel et al. distinguished graphics via four quadrants based on dimensions of the reference grid (e.g., Cartesian coordinates, or a map) and the information (e.g., points on a scatter plot, or 3D bars): 2D reference with 2D data, 2D reference with 3D data, 3D reference with 2D data, and 3D

[47] Büschel et al., "Augmented Reality Graph Visualizations: Investigation of Visual Styles in Three-Dimensional Node-Link Diagrams."

[48] Büschel et al., "Augmented Reality Graph Visualizations"

reference with 3D data (fig. 5).[49] With this categorization the authors were able to identify fundamental differences among existing visualization techniques, in an effort to streamline decision-making. Because each category had its own benefits and drawbacks, the authors concluded it was not possible to select one category that was best for every situation.

Because augmented reality is an emerging field, there has no substantive advancement of Dübel et al.'s classifications or theoretical frameworks as they relate to augmented reality visualization. That is in part because there are few real-world applications at the moment.

The rise of news graphics and innovation in news

One of the widest-used applications of augmented reality, including data visualization, is within the news industry. Beginning in 2013, a number of newsrooms experimented with storytelling using mobile-app-based AR. The first examples applied AR on top of a print newspaper, pulling up a webpage when a user pointed their camera at a headline or overlaying a video over a static photo in the paper.[50] But technology at the time didn't support full-blown 3D graphics, which emerged with the 2017 and 2018 release of Apple's ARKit and Android's ARCore, respectively. Once AR was widely available in smartphones, multiple newsrooms began to experiment with augmented reality. These efforts typically originated with existing graphics teams, who are often some of the first sources of new storytelling techniques in newsrooms.

News graphic teams, of course, have existed for decades, responsible for designing charts, infographics, and illustrations to add different forms of storytelling to print newspapers.

---

[49] Dübel et al., "2D and 3D Presentation of Spatial Data: A Systematic Review."

[50] Pavlik and Bridges, "The Emergence of Augmented Reality (AR) as a Storytelling Medium in Journalism"

One amateur historian of the field, John Grimwade, a former graphics editor at *The Times of London* and *Conde Nast Traveler*, points to early news graphics as far back as 1804, although the field grew dramatically in the mid-twentieth century, particularly in newsmagazines like *TIME.*[51] As news organizations' websites became increasingly important to the reader experience in the 1990s and early 2000s, some outlets experimented with more interactive, explorable story formats that required a hybrid of computer science and journalism skills to produce, which many newsrooms began to call "interactive journalism."[52]

For years now, these hybrid teams of journalist-developers have operated at the cutting edge of news, often creating new formats and testing out new technologies within a newsroom. As the teams have developed, some newsrooms have launched even more formal innovation units like R&D teams at the *New York Times*[53] or *Wall Street Journal*,[54] or the "Advanced Story Lab" at the *Washington Post*.

All of these aforementioned teams have experimented with AR in different capacities, usually phone-based projects that are accessed through the newspapers' mobile apps and made possible by Apple's release of ARKit in early 2017.[55] Many of these stories were pegged to scheduled news events like the Olympics or New York Fashion Week, in part because the AR development process is still relatively difficult. However, the New York Times, in particular, has

---

[51] Usher, *Interactive Journalism,* 39-40.

[52] Usher, *Interactive Journalism*, 18.

[53] "The New York Times - Research & Development." https://rd.nytimes.com

[54] Nunes, "WSJ R&D."

[55] See Appendix F for a partial list of AR stories published by major US news organizations

improved their speed for AR projects, publishing AR stories for breaking-news events like a group of young boys stuck in a Thai cave,[56] or a fire in the Notre Dame cathedral in Paris.[57] The *Times*,[58] the *Post*, and others have also been improving their capabilities to do "photogrammetry," using hundreds of still photographs to reconstruct an object in 3D for AR or VR. In a July 2019 talk in Brooklyn, the (now-former) head of "Immersive Storytelling" at the *Times*, Graham Roberts, said that photogrammetry feels "inherently journalistic" since it was so similar to photography in its representation of real life, and the process is becoming increasingly easy to do on-deadline or in the field. As a result, many of the early AR stories have relied on this technique, instead of using fully-animated or abstract data visualizations.

A few newsrooms have experimented with data visualizations in AR, including a daily chart of market performance from the *Wall Street Journal*,[59] a 3D explanation of data from CERN's Large Hadron Collider in the *New York Times*,[60] and a *Times* project that superimposed polar winds on a 3D globe.[61] However, these seem to be the exception, rather than the rule, when it comes to AR coverage by these major news outlets.

---

[56] "Step Inside the Thai Cave in Augmented Reality." July 21, 2018. (https://www.nytimes.com/interactive/2018/07/21/world/asia/thai-cave-rescue-ar-ul.html)

[57] "Why Notre-Dame was a Tinderbox." April 17, 2019. (https://www.nytimes.com/interactive/2019/04/17/world/europe/notre-dame-cathedral-fire-spread.html)

[58] Porter, "Documenting the World in 3D."

[59] "Experience augmented reality in the WSJ app for iOS." September 21, 2017. (https://www.dowjones.com/press-room/experience-augmented-reality-wsj-app-ios/)

[60] "It's Intermission for the Large Hadron Collider." December 21, 2018. (https://www.nytimes.com/interactive/2018/12/21/science/cern-large-hadron-collider-ar-ul.html)

[61] "A Closer Look at the Polar Vortex's Dangerously Cold Winds." January 30, 2019. (https://www.nytimes.com/interactive/2019/01/30/science/polar-vortex-extreme-cold.html)

While many of the AR-driven stories received praise from the industry, their creators have admitted there's still much work to be done. Roberts explained how his team discovered that users still didn't engage with AR to the degree that they expected, and came to the conclusion that it wasn't often worth the effort put into development.[62] The *Times*, he said, has since moved away from AR and towards in-browser 3D experiences, having concluded that mobile AR is still an awkward format and readers don't always want to spend the time or effort interacting with it. He noted that the *Times* still uses AR features for some stories (like a recent feature for the Apollo 11 moon landing anniversary[63]) but usually as an optional add-on experience, as opposed to the default.

Still, Roberts expressed hope that future AR technologies would open up new opportunities for AR, since it clearly holds storytelling potential — just not always on mobile devices. That's why news organizations continue to invest in the tools and practice of AR story development: because, as Roberts explained, they hope they be ready for newer technologies as they become available.[64]

---

62 Roberts, "Augmenting Journalism: Immersive Storytelling at The New York Times."

63 Roberts and Corum, "How We Augmented Our Original Reporting of the Moon Landing for Its 50th Anniversary."

64 Roberts, "Augmenting Journalism: Immersive Storytelling at The New York Times."

Methods

To channel the advantages of both qualitative and quantitative studies, this study used a mixed-methods approach to better understand the best practices for data visualization in augmented reality. The first portion of this study applied the 2D design theories from Tufte, Bertin, and others to the new design context of 3D space. The process of design is important in itself: while it may be somewhat subjective, it replicates the development process of graphics teams in many newsrooms, which don't often test their graphics with audience members before publication.[65] Then, a second portion used user testing to understand how readers interact and interpret the AR graphics created during the first section.

Design Process and Theory

This thesis required the creation of new tools and technology to produce 3D charts, as well as an iterative design process (as described by Tufte[66]) to improve the clarity and design style of the graphics through repeated re-designing. Ultimately, this section of the research produced *guidelines* for designing better visualizations in AR — it was not intended to produce hard-and-fast empirical rules.

A single topic was chosen for all three graphics — campaign finance in the 2020 Democratic primary — to give the project a unified theme and better simulate the experience of reading AR visualizations as part of a news package. The three graphics were composed inside

---

[65] Usher, *Interactive News,* 170.

[66] Tufte, *Visual Display*, 123.

one "tappable" article format (an increasingly common one for stories on mobile devices[67]) and randomized in order to avoid recency bias in the following survey results.

To test the simplest approach to 3D graphics — placing a very traditional graphic form in an extra dimension — one of the charts in the story was a relatively simple bar chart, showing the total amount of money each candidate raised. This bar chart was augmented with a second dimension, with bars showing poll numbers for each candidate pointed in the reverse Z axis away from the user. This further simulates a naive approach to 3D graphics, which uses the extra dimension of space to quickly and easily add an extra dimension of data.

Another graphic showed a map of donations by county, colored by the top candidate who donated to that location. This graphic attempted to use the extra dimension to make geographic data clearer — since 2D maps are often constrained in the detail of data they can contain, they have the most to gain from an extra dimension of space to work with.

Finally, the third graphic showed a collection of abstract spheres, scaled and colored by the size of different types of donations to presidential campaigns. Each candidate had a green sphere for small-dollar donors (less than or equal to $200), a yellow sphere for large-dollar donors (greater than $200), and a blue sphere for self-funded campaign donations from the candidate. This graphic attempted to use all three dimensions to show relative scale in a way that simulates how one might compare the size of objects on a table in real-life.

[67] e.g. "China, Emergence of a Trade Leviathan," *The Wall Street Journal.* (https://www.wsj.com/graphics/china-emergence-of-a-trade-leviathan/)

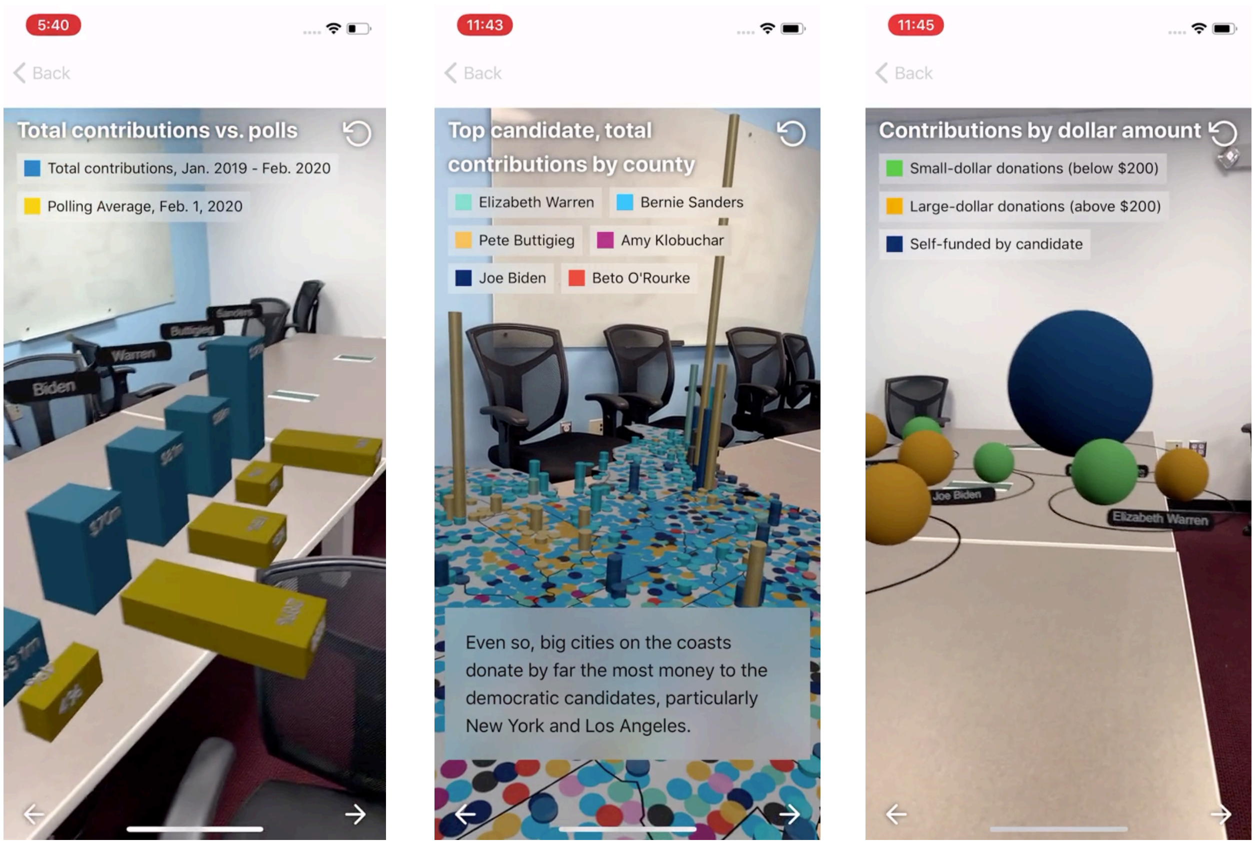

Figure 6: Views of the bar chart (from behind), the map, and the spheres graphic, from screenshots taken during a user test.

All three graphics were initially designed and created using a custom tool built on the Three.js javascript library,[68] which were then exported to a GlTF file and converted into Apple's preferred USDZ format. This made it relatively simple to display the graphics using Apple's RealityKit framework[69] inside an iOS app, or inside Apple's QuickLook AR feature[70].

For this test, an AR-enabled news app previously developed by the researcher within UNC's Reese Innovation Lab was used to display the AR graphics, making it simple for the researcher to design, develop, and publish an article on the app for this study. However, the

68 "three.js - JavaScript 3D library" (https://threejs.org)

69 "RealityKit | Apple Developer Documentation" (https://developer.apple.com/documentation/realitykit)

70 "QuickLook | Apple Developer Documentation" (https://developer.apple.com/documentation/quicklook)

technology used is not directly relevant to the experience being tested, since this study focuses on the specific design aspects of the graphics *within* AR, which can be reproduced on multiple different apps or AR platforms.

User Testing

Once they were designed, the graphics were evaluated by users in order to get a better, more empirical understanding of which design aspects worked well and which were confusing. A sample of 23 undergraduate participants were drawn from the UNC Hussman School of Journalism and Media's research pool, a sufficient size for both the qualitative[71] and quantitative[72] aspects of the study, as shown by the work of user experience experts Jakob Nielsen and Don Norman. This research was governed by UNC's Office of Human Research Ethics, and after review the following methodologies were determined to be exempt under regulations for Benign Behavioral Interventions.[73]

After signing a consent form (Appendix A), participants completed a brief background survey (Appendix B). This survey collected basic information about their news reading habits and understanding of augmented reality, to allow comparisons between highly news- and tech-literate users, and others who may have less familiarity with the two.

Then, the users were given a device pre-loaded with the AR experience (Appendix C); approximately half (11 participants) on an iPhone X, and half (12) on a 12.9" iPad Pro. This variation in devices is also important for understanding real-world usage of AR, where news graphics are consumed on a wide variety of device types and screen sizes.

---

[71] Nielsen, "Why You Only Need to Test with 5 Users"

[72] Nielsen, "Quantitative Studies: How Many Users to Test?"

[73] UNC Study #20-0034

Throughout the user test, the researcher asked the participants a few open-ended questions following the user testing methodologies set out by the Nielsen-Norman group.[74] That methodology includes encouraging users to verbally explain their thought process during their interactions with the AR models, along with careful observations of their behavior.[75] As Freitas, et al. noted in their 2014 review, user experience methods can be useful tools for studying information visualizations, offering an alternative understanding of the user experience oriented around their own goals and intuitions.[76] This form of research generates insights into what aspects users found most interesting, what parts were most confusing, and how well the user interface matches up to users' instinctual expectations. Ultimately, these qualitative observations provided exploration and elaboration on some of the study's quantitative aspects, and produced useful insights that could allow for the improvement of the graphical design in this study.

Drawing upon the graphical perception research methods developed by Cleveland and McGill,[77] the researcher then provided each participant with a brief survey to measure their comprehension and retention of information in the charts (see Appendix D). These questions reflected the goals of the average news consumer, seeing whether they understood and recalled the main points of the graphics that accompany the story. This survey technique follows the human-centered visualization theory developed by Frietas et al. and Shah et al. In addition to comprehension questions, the participants were presented with questions to rate on a 1-to-5

[74] Nielsen, "Usability 101: Introduction to Usability"

[75] Nielsen, "First Rule of Usability? Don't Listen to Users."

[76] Frietas et al., "User-Centered Evaluation of Information Visualization Techniques: Making the HCI-InfoVis Connection Explicit."

[77] Cleveland and McGill, "Graphical Perception"

Likert scale their impression of different aspects of the experience; which was modeled on Yang et al.'s study of globes in virtual reality.[78]

Finally, participants were interviewed with a series of in-person prompts, as well as follow-up questions to review specific comments and clarify parts of the overall experience. (See Appendix E for a list of the standard in-person interview questions.)

All of the results were be applied toward a few research questions:

1. Are users with previous experience in augmented reality more likely to understand the contents of these augmented reality visualizations?
2. Are certain types of graphics better suited to 3D environments than others?
3. Do users find 3D graphics more or less appealing than 2D ones?

In the spirit of most usability testing, the questions are only guidelines for the research: the primary goal of the study is to yield unanticipated insights into the design of AR visualizations. The purpose of this study is not to prove or disprove a hypothesis, but instead to gather a set of observations and best practices that will guide the development of data visualizations in augmented reality in the future.

[78] Yang et al., "Maps and Globes in Virtual Reality."

Results & Discussion

The initial intuitions and testing during the design process closely mirrors the process of constructing graphics in the newsroom, so observations from that phase are recorded here. Then, the examination of the results of users' test experiences can further validate or refute those design intuitions, helping form a stronger set of guidelines going forward.

Design Process - Bar Chart

The AR bar chart was initially designed to closely emulate a traditional 2D bar chart, so the early design process was straightforward. In its initial conception, it would serve as a simple way to highlight the disparity in fundraising between different candidates. While constructing the chart in the editor (on a normal computer screen), it initially felt difficult to interpret the heights of different when viewed at an angle — as discussed in the literature review, a number of earlier visual perception studies have shown that 3D perspective can distort users' interpretation of relative height[79].

Interestingly, that concern didn't translate to the 3D space: it felt surprisingly natural to view and compare the sizes of each bar when the designs were tested in augmented reality on a mobile device. The ability to move around and change perspective largely felt like it mitigated the distorting effects, making it easier to compare values accurately.

Another crucial design element are the labels, which are equally important to helping users interpret the graphic as the bars themselves. Initial iterations of the design included a Y axis to represent scale in height, but the scale bar felt out of place and hard to use. Since every bar in this 3D chart felt like a discrete object, it felt more natural to directly label each bar's value on its

[79] Lind et al., "Metric 3D structure in visualizations," 54; Bleisch, "Toward Appropriate Representations of Quantitative Data in Virtual Environments."

own surface. However, this reveals new problems: because labels are two-dimensional, they can only face one direction on the bar chart. To make the labels more visible from all directions, they were drawn on both sides of each bar, which partially solved the issue, although the labels are still a little difficult to read on the smallest bars.

Another solution to the labeling issue came from Apple's new Reality Composer software,[80] which allows objects to freely rotate to maintain their angle towards the user. This was not as helpful for the value labels (which could not freely rotate without being occluded) but was a useful way to label the names of each candidate above their bar (fig. 7).

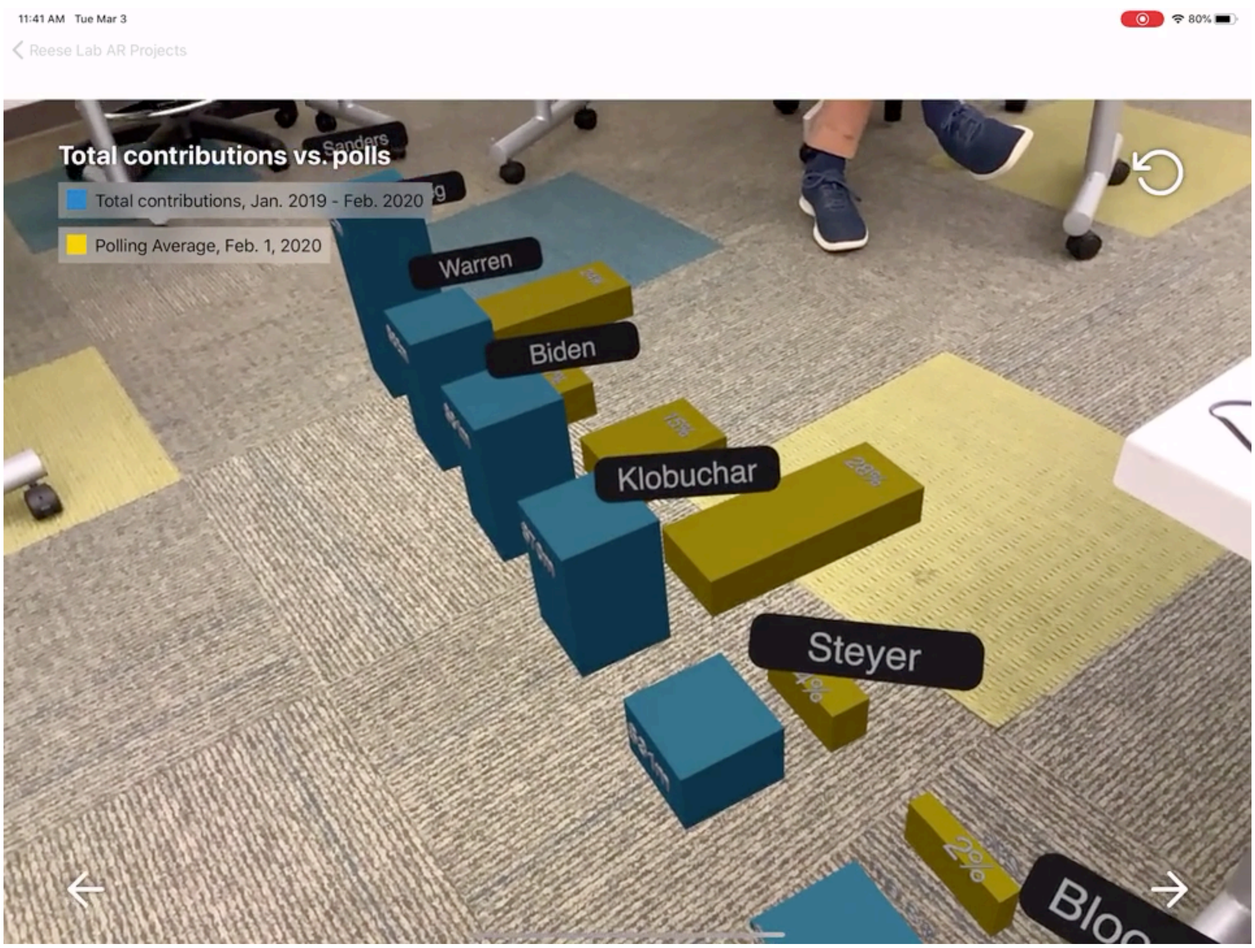

Figure 7: Bar chart labels which automatically rotate to face the viewer, seen in a screenshot from a user test.

[80] "Augmented Reality – Reality Composer" (https://developer.apple.com/augmented-reality/reality-composer/)

Once the candidate names were visible from all angles, adding a second dimension of data (candidate poll numbers) became possible. Early iterations of the design located the second set in front of the upright bars, pointing towards the user, but that arrangement obscured the height of the upright bars and made them difficult to judge. Since the fundraising was still the primary variable shown in the chart, it didn't feel right to obscure them behind other bars.

Eventually, the poll number bars were placed in the back. It came with a downside: sometimes those bars would be obscured if the user viewed them at a low angle, but if the user looked from a higher angle (e.g. by standing up) they would be more visible. Additionally, this placement would test user behavior: observing how users explore the second set of bars provides data about how people move around to get into a more ideal view of the graphic.

Design Process - US Map

AR offers many advantages in designing maps, which are already difficult to display effectively in 2D. Geography requires an X and Y dimension, so any data layered on top of it in traditional graphics must be carefully designed to stand out.[81] Despite this challenge, maps also offer a high amount of *data density* — information encoded in a relatively small space — so a 3D translation is all the more valuable for readers to learn a lot from a relatively small graphic.

Many news organizations have published highly detailed maps of donations to 2020 primary candidates,[82] so this AR map was constructed to add an additional variable of scale on top of the already colorful 2D maps out there. The implementation was relatively simple, by drawing a cylinder with height representing donations from each county, colored by the

---

[81] Tufte, *Envisioning Information*, 28

[82] "Detailed Maps of the Donors Powering the 2020 Democratic Campaigns," *The New York Times*, Aug. 2, 2019. (https://www.nytimes.com/interactive/2019/08/02/us/politics/2020-democratic-fundraising.html)

candidate who got the majority from that location (fig. 8). Surprisingly, it worked well from almost the first conception, because the density and variety in height of the cylinders made the data feel more alive and explorable. Tufte recognized this trait of good graphics, calling it "micro/macro" composition: adding so many details to a visual that the viewer is simultaneously drawn into small details and the big picture.[83] The map allowed exploration of certain familiar places— maybe hometowns, or favorite cities — along with big-picture trends like Bernie Sanders' widespread support and highly concentrated high-dollar regions for candidates like Kamala Harris and Pete Buttigieg.

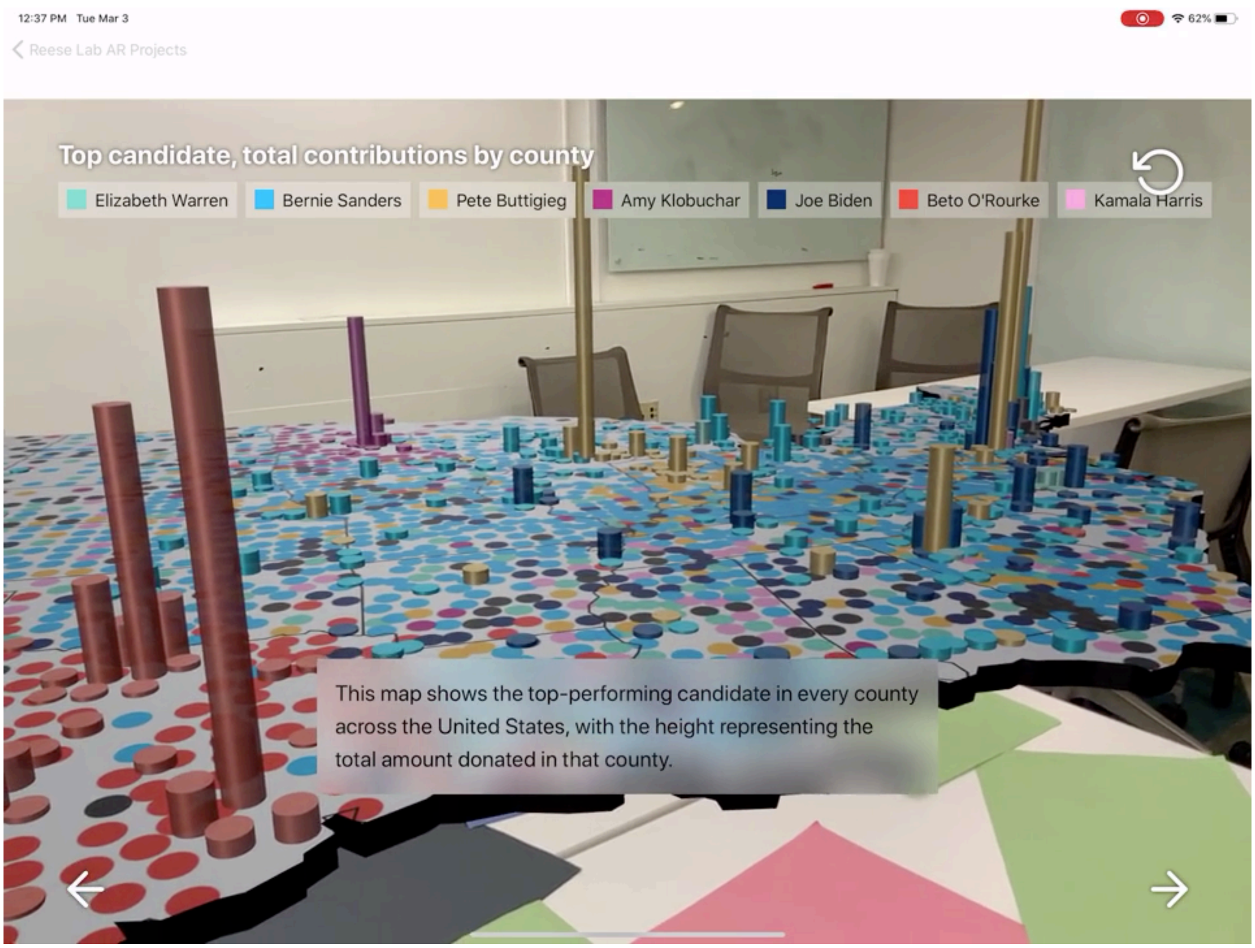

Figure 8: View of the map in a screenshot from a user test. (Colored paper was placed on the table to improve tracking.)

[83] Tufte, *Envisioning Information,* 37.

The AR graphic did require a fair amount of fine-tuning for performance, because the thousands of vertical bars tended to slow down rendering speeds and degrade performance on the device. But after some adjustments that flattened nearly flat cylinders (less than one millimeter) onto the surface itself, the number of shapes was reduced enough for devices to render the graphic efficiently.

One of the most complicated design challenges for this graphic was labeling. Initial iterations included direct, colored labels in 3D space near where each candidate had a cluster of donors, which were useful when the user took a step back to view the whole graphic but difficult to use when referring back and forth between different regions. Ultimately, a fixed key with color swatches and candidate names was added to the 2D screen (not the 3D asset), where it was easier to glance at as a reference tool. This suggests that AR development might still take advantage of 2D screen space, in certain cases where reference or permanence is useful. The key is to balance 2D context (like scales or keys) with visible space, so that on-screen content doesn't crowd out the window of AR content.

Design Process - Spheres

The last graphic uses spheres to compare the composition of donations to specific campaigns, somewhat inspired by a 2D graphic from FiveThirtyEight that used bubble shapes to show big donors, small donors, and self-funded campaign contributions.[84] After experimenting with different grid layouts and orientations for spheres, the clearest appeared to be a grouping method: enclosing various component spheres for a given campaign within a surrounding circle for each campaign (fig. 9). This layout felt the most natural because it simulated the placement

[84] "What Third-Quarter Fundraising Can Tell Us About 2020," *FiveThirtyEight*, Oct. 16, 2019. (https://projects.fivethirtyeight.com/2020-fundraising-q3/)

and arrangement of real objects in space, as if one sat down at a table and grouped a set of marbles in different ways.

Critics of the design might point out that volume is hard to measure and compare, and they would be right: the graphic made it difficult to point to a given sphere and read out a precise value for that contribution type. Initial renders of the design on the 2D computer screen prompted some of the same concerns, but once viewed in 3D it felt notably clearer and easier to compare. The main point of this graphic wasn't the details: it was designed to highlight a more simple contrast between campaigns' fundraising strategies. Since the other graphics in this study

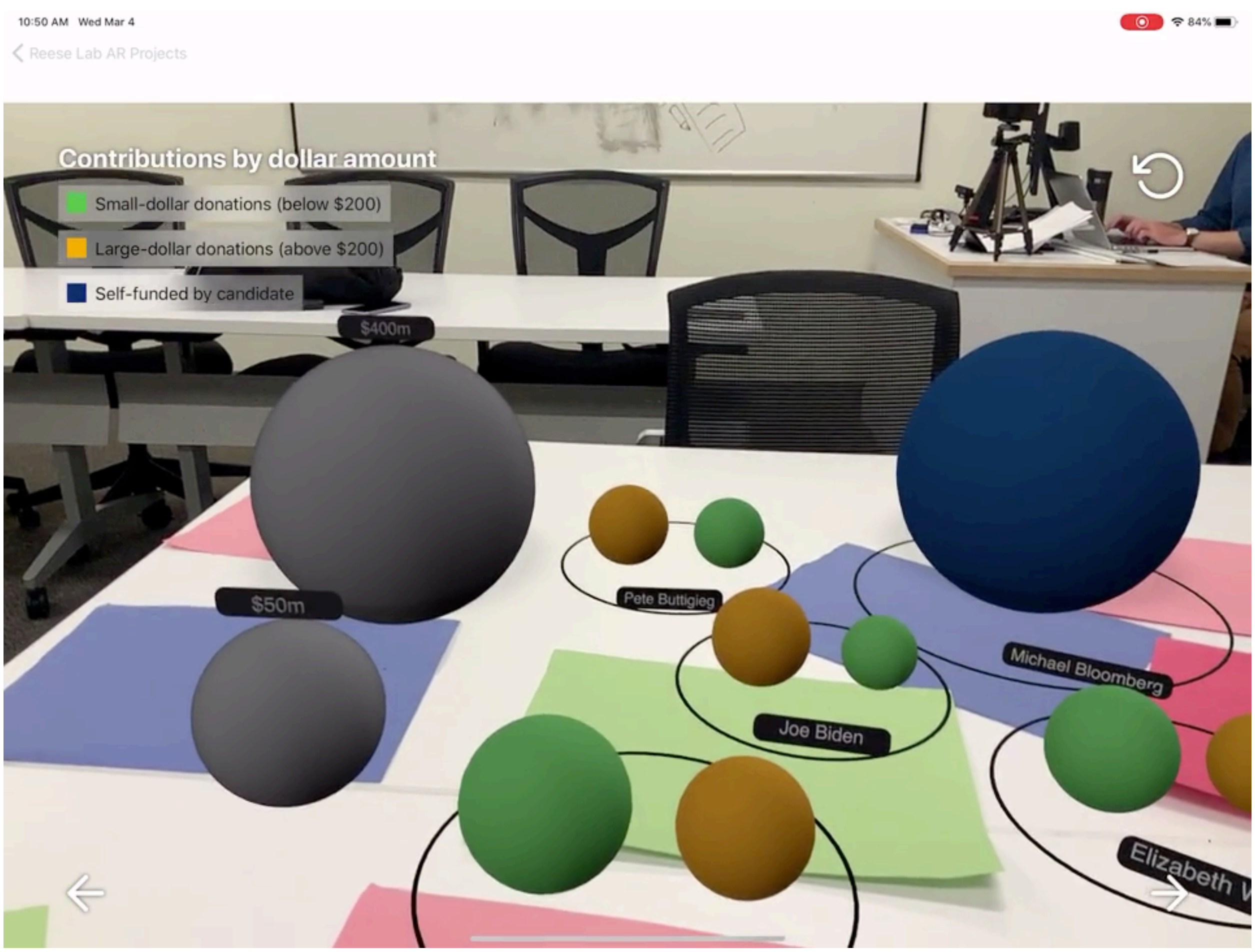

Figure 9: The spheres graphic seen in a screenshot from a user test, with Michael Bloomberg's large self-funded sphere on the right, and grey scale spheres on the left. (Colored paper was placed on the table to improve AR tracking.)

are much more numeric and precise, it seemed relevant to examine a simple, but hopefully impactful, form of design.

The most obvious takeaway from the graphic was the massive blue sphere showing Mike Bloomberg's colossal, self-funded campaign coffers, which stood out from the others by an order of magnitude in scale. It was also easy to tell whether a campaign relied more on small donors or large donors based on which sphere was larger, making it possible for readers to draw conclusions about campaign strategies.

Labels were still a bit tricky (an apparent trend for 3D visualization design): the eventual design included a color key affixed to the 2D screen as in the map, and direct, auto-rotating labels for each campaign as in the bar chart. Scales were even more difficult, since there was no clear way to direct-label each sphere (without cluttering the design) or include a reference scale in 2D (since a sphere's relative size versus the screen depends on how close the viewer is). The ultimate solution was to add a set of grey reference spheres for scale off to the left side, which partially solved the problem but still made it difficult to measure the exact size of a given sphere.

User Testing Observations

Even after all the improvements generated during the iterative design process, user feedback opened up a number of new (and occasionally surprising) insights into designing AR experiences. Every designer's own biases, experience, and habits will inevitably lead to blind spots in their designs, which user testing can help uncover and improve.

The first interesting challenge users encountered was *placing* the objects. While the UI for placing AR objects around the room wasn't strictly part of the design process for this thesis, it undeniably contributes to a user's experience, either positive or negative. The placement interface mimicked ones used by the New York Times and Apple, where a shadow moves around the room wherever the device is pointed, and the user can tap to place the object where the shadow rests at any time. Users' comfort with this process can be measured by the total time it takes them to place objects; a shorter time implies a better user experience that gets viewers to the AR content faster.

Test subjects completed placing their first object in 26 seconds on average, although there was a slight difference between iPhone users (22.7 seconds) and iPad users (29.8 seconds), which may be due to the iPhone's better tracking capabilities. Users who had prior experience with AR (measured based on a correct definition for "augmented reality"[85]) were the fastest to place objects, at 20.4 seconds vs 29.7 seconds.

Table 1: Average time for users to place items in AR

	Device type		AR experience		Overall
	iPhone	iPad	Yes	No	
Avg. placement time (s)	22.73	29.83	20.38	29.67	**26.43**
N (users)	11	12	8	15	**23**

When asked after the experience, most users said the object placement interface made sense to them. But a few said they had difficulty placing it in an ideal location: one user placed a chart just a few inches in front of them, then said it felt uncomfortably close. Another said "the first time I didn't really understand how to place [the charts], but after that it was easier to place." Overall, 8 of the 23 users mentioned the placement process when asked when they felt

[85] This measurement was chosen based on comparison of user answers to "Have you used AR before?" and "Define Augmented Reality." It turned out that some users who said they had used AR were unable to define it, introducing a possible error that could skew the results. It seems safer to include users who understood AR but never used it in the category of "AR experienced," as opposed to users who *said* they used AR but couldn't actually define it. As a result, a users' correct definition for "Augmented Reality" will be used for a proxy for their AR experience throughout the rest of the paper.

most confused. Those experiences, combined with the data, suggest that placing objects in AR is a learned skill — so it's crucial to provide users with the ability to try multiple times with a "do-over." The app included a reset button in the top right corner, but only a few users took advantage of it, while others simply chose to cope with the slightly uncomfortable placement location. The users who used the reset button said they appreciated it, explaining that the ability to retry made them feel more comfortable and ultimately helped them place the graphic in a position easier to view.

Many users also remarked that they were reluctant to move around. One said, "I wouldn't think to get up and look around more closely." Another explained, "Normally when you're using an article you don't think about the physical space you're in." Despite that walking around a scene to view it from different angles is a key part of AR, the feature wasn't intuitive to the users. Even those who moved around recognized that they might not in other contexts: "A lazy person wouldn't want to get up and see it," one user said. "I would rather move on instead of stopping to do AR in certain situations."

Yet when users remained seated it opened up other usability and design issues. For example, a number of the seated users didn't notice the second set of bars showing poll numbers in the bar chart, because they were occluded behind the front fundraising bars when viewed from a low angle above the table. Others noticed them, but felt they were difficult to evaluate. "I kind of wish the yellow bars were next to the blue bars," one user said. "I find it difficult to compare the differences between the percent and the money amount from this angle."

In addition to occlusion, the fact that users didn't know to move around meant that off-screen items were confusing to them. "At first, I was only keeping it level here," one participant said, gesturing at the table while looking at the bar chart. “But I had to move it up before I saw the labels over each bar.” Without any sort of hints, users (especially those unfamiliar with AR) sometimes didn’t notice entire components like direct labels or scales.

Eventually, however, most users stood up and moved around. “At first I had a panoramic instinct, I thought it was only going to [surround] where I was sitting,” one explained. “I thought I had to stay in the same place, but if I move around it I can see much better.”

The numbers appear to support that trend: more than half of the users eventually moved around, even if they expressed hesitancy at first. Eight users moved around on their own, and an additional five looked confused and asked permission before standing up. Six others moved or rotated in their spinning chair, but never stood up. Four users didn’t move at all, opting to remain seated.

Table 2: Breakdown of movement types

	Device type		AR experience		Overall
	iPhone	**iPad**	**Yes**	**No**	
Stood up, moved around	45%	25%	63%	20%	35%
Moved, but asked first	18%	25%	0%	33%	22%
Moved in chair	18%	33%	25%	27%	26%
Did not move	18%	17%	13%	20%	17%
N (Users)	11	12	8	15	**23**

As with placement, user movement tended to correlate with a prior understanding of AR. Of users who were correctly able to define “augmented reality,” 63% (5 of 8) walked around the room on their own, without asking for permission or help. For users who couldn’t define the term, only 20% (3 of 15) moved around on their own. More iPhone users moved than iPad users

(45% vs 25%), suggesting that the phone's portability encouraged may have encouraged people to move more.

Beyond these overarching trends, a few observations specific to each graphic came up frequently during the user tests.

User Testing Observations - Bar Chart

The most common complaint about the bar chart, by far, was that it was difficult to see the horizontal poll number bars in the back unless the viewer was standing up or walking around. That was a complaint listed by 13 of the 23 (56%) of the participants who were asked what parts of the entire article most surprised or confused them, far more than any other feature. A number of users also pointed out that they felt the graphic would be just as effective and maybe less confusing in 2D, so the AR felt unnecessary for displaying the data in these bars.

When asked which graphic was their favorite, only 2 of 22 respondents (9%) said the bar chart felt clearest to them. The low rating might be because they felt it was difficult to understand: only 43% of respondents could correctly identify who raised the most money in the survey, and only 19% of respondents correctly chose which candidate was polling highest. The result implies that users' confusion wasn't just a feeling: it had actual impacts to their comprehension and recall.

User Testing Observations - US Map

The map was the most consistently well-liked of the graphics, with 64% of participants naming it when asked which graphic was clearest. "Everyone knows the US map, so it's easy to comprehend what it's trying to get across," one participant explained.

This familiarity was a common theme among user responses. “It’s a little personal,” another explained. “You’re from a certain state, and you want to see how much your state gave, or the surrounding states. Or [you think] ‘Oh, my friend’s from here… so I want to see how it looks over there.’”

Interestingly, users also thought the map was complex: that description was mentioned by six of the fourteen participants who wrote an explanation of why the map was the most interesting. This suggests that complexity isn’t inherently the enemy of clarity: as in Tufte’s theory of micro/macro composition, complexity sometimes makes for more engaging graphics. And after viewing it, readers took away some useful facts, since 65% of participants were able to name a specific detail they learned from the map (as opposed to a more vague statement). This was especially true for users who had used AR previously (86% vs 57%), suggesting that users more comfortable in AR might have paid more attention and recalled more specific details from the map.

Table 3: Percent of users reporting specific learnings from the map

	Device type		**AR experience**		**Overall**
	iPhone	**iPad**	**Yes**	**No**	
Recalled specifics	64%	67%	88%	53%	65%
N (Users)	11	12	8	15	**23**

As with every graphic, the map also had opportunities for improvement. A number of users pointed out that the large number of colors in the key made it difficult to view differences between candidates. This was further impacted by the way iOS’s “physically based rendering” shows colors in AR: when trying to look realistic, it simulated the colors and lighting in the environment, which made the AR colors deviate slightly from the color of the key. This made it difficult for users to uniquely identify the candidate for specific bars, especially those of similar colors (e.g. pink and purple, or green and blue).

User Testing Observations - Spheres

The spheres graphic was the most controversial one: users seemed to either love it or hate it. A number remarked on its simplicity, calling it "easy to understand," or "simple to compare" when describing it as the clearest graphic. When asked to report one fact they learned from the whole article, 61% (14/23) mentioned something about the small and large dollar donors, or more frequently about Bloomberg's self-funded campaign. The result suggests that this main point[86] was communicated clearly, since it stuck with so many users after the fact.

However, other users disliked the graphic, complaining that it felt imprecise and hard to evaluate the scale of each sphere. Some didn't notice the grey size scale on the left side, and others thought the clustered layout felt hard to compare. "The spheres felt too abstract, I wish there was more concrete detail," one user noted. Another said, "it was sort of hard to distinguish between the spheres, but I could see the frontrunners and stuff."

Those who noticed the scale also expressed dissatisfaction, which was they said hard to interpret and not clearly labeled enough because the gray spheres were so far to the left. While a few immediately recognized that they represented a scale, most did not — and either asked the researcher what they meant, or simply moved on without any explanation.

Both sides of this dichotomy are expressed in the comprehension question answers. 70% of participants correctly answered a question about the big takeaway (that Michael Bloomberg self-funded his campaign and had a huge budget), but only 48% were able to correctly identify

[86] See discussion in "Design Process - Spheres" section above.

the cut-off between “large” and “small” donors mentioned in the story and on the graphic key. (The answer was $200).

This might show that there are two different types of users: some who want details, and others who just want the broad picture and one point to remember. An effective graphic would embrace both viewings, offering macro themes and micro details to users who want each.

Other Observations

When the methodology was designed, the charts were shuffled for each user to decrease recency bias — avoiding a situation where users might have more accurate memories (and therefore answers) about the graph they most recently seen. However, that appeared not to matter: there was no consistent pattern between visualization order and accurate answers to the survey. For the bar chart questions, participants who viewed the chart most recently actually performed the *worst* and those who viewed it first did better (15% vs 38%). Participants who viewed the map last answered its questions slightly better (69% vs 57%), but participants who viewed the spheres last did worse (48% vs 73%). These conflicting responses imply that there was no consistent pattern: the differences could easily be due to random chance, given the small sample sizes.

Table 4: Percent of correct answers for each graphic, based on order of appearance.

	Bar Chart		**US Map**		**Spheres**		**Overall**
Questions:	First	Last	First	Last	First	Last	
Bar Chart Qs	38%	15%	13%	38%	40%	38%	31%
US Map Qs	36%	58%	57%	69%	76%	42%	57%
Spheres Qs	58%	46%	42%	79%	73%	48%	59%

Note: Only includes non-binary multiple choice questions, and excludes “trick” questions that less clear correct answers and very low accuracy rates.

It is also notable that so many of the multiple choice questions had low correct rates. That might suggest a lack of attention paid to details in the graphics, especially since participants were not told they would be quizzed afterwards. While some of the pairwise comparison questions were answered more accurately, they were also more forgiving to guessers and more likely to be affected by outside political knowledge. As a result, those pairwise questions were excluded from this analysis, to keep the results more focused on similar styles of questions.

Even with all their complaints and possible improvements, users overwhelmingly indicated that they would enjoy more AR data visualizations: 19 of the 23 (83%) of participants said they would use the technology to read other news and information articles if they were available. Most noted the caveat that they would only use AR when they had enough time, and likely at home and not in public. Because the test pool was undergraduate students, many thought about applying the technology to their classes, surmising that they would better remember facts from textbooks or readings if they were forced to explore them in augmented reality.

Yet, other students recognized the challenges in more widespread adoption of AR in news articles like the study. “Augmented reality is fighting with tradition,” one said. “We’ve been doing this for so long, that to insert this new graphics system, you’re fighting with people on how they read news.”

Conclusions

The results of the user testing discussed above have been integrated into a few recommendations for designers looking to develop augmented reality data visualizations that are clear, engaging, and easy to use.

Design Recommendations

1. **Try to keep all items in view at the start.** This solves many of the problems where users felt like they didn't notice a specific detail (e.g. a label) because it was above or to the side their device screen. This is especially important when users are placing large objects close to themselves, because it is confusing to feel as if you have been surrounded and don't know where to look. This also applies to occluded objects, like the horizontal poll numbers in the bar chart graphic in this study. If objects must be placed behind another, try to ensure that the viewer's initial angle is such that they can see at least a piece of what they are missing. Or, in a worst case scenario, use a system of highlights or pointer arrows to help users learn which way to direct their attention.

   Another possible approach is that of Apple's QuickLook AR software: when an object is placed into AR, it "grows" into place a using scaling-up animation so that users get a glimpse of the full scope of the graphic. This allows them to build a mental map of the space, even if they can't see it all at once.

2. **Optimize for "micro/macro" composition. (Make a big point *and* allow for exploration.)** Like Tufte's description of good 2D graphics, the best AR graphics offer one key takeaway users can understand quickly, but also provide explorable details so they can dive in deeper on their own. The map in this study appeared to do this well, but the bar chart and the

spheres were less successful. Both are essential: without a big point, users might feel lost and confused; without a deeper level to explore, users might feel unsatisfied with the result. Since AR takes so much time and effort to set up, you don't want a user to feel like they've wasted their time on something they could've learned through a much quicker 2D chart. On the other hand, AR gives users a sense of agency to be able to explore a visualization — good designs should offer up those details and invite the user in.

3. **Assume users won't move around.** Again and again, participants explained that even though they stood up and walked around during the study, they were unlikely to do so if they were casually scrolling through news articles on their own. Maybe the most passionate viewers will really engage with a scene, but most will stay sitting where they were and merely shift their phone camera to look in different directions. As a result, AR experiences should be designed in such a way that all the information is the right scale and position to be viewed from close by, without assuming that users will spend time exploring every corner.
4. **Use labels sparingly, and make sure they rotate to follow the user.** Labels are an essential part of any graphic, helping tell the user exactly what they're looking at. This becomes significantly more difficult in AR, where you can't make assumptions about what is in or out of view. Furthermore, occlusion from other objects in the scene and the ability to view from multiple angles means that text is sometimes illegible, even if the label itself is visible. Therefore, labels are best used only when necessary and should be located in a way that is constantly visible, like above the entire scene, instead of in the middle of everything.
5. **Allow opportunities for users to learn and fail.** One of the recurring themes in this research shows that users who had previously used AR before this study were able to load

graphics faster, move around more, and retain more from the graphics after they finished. AR designers shouldn't assume everyone is an expert, but they should provide tools for users to get situated and comfortable in AR environments before dropping them into a more complex experience. That includes clear and frequent textual instructions, and especially tools that help with placement. A number of users in this study were delighted to find the "reset" button that let them re-place graphics in a different location, saying it gave them a sense of control and helped them correct early mistakes before they understood what would happen. This improved the user experience, helping them learn faster and better enjoy the visualizations themselves.

While these recommendations are not a comprehensive list, they can form the basis for a further discussion of best practices for data visualizations in augmented reality. Future research could explore different forms of augmented reality (e.g. glasses and goggles versus smartphones), different interaction patterns (by creating interactive visualizations unlike the static ones used in this study), or scale a similar methodology to a larger sample size in order to better tease out differences with higher statistical certainty.

As augmented reality devices get better and better, it's only a matter of time until a larger number of consumers want to consume news content and data visualizations in AR. By studying best practices and improving designs today, we can improve the value and usage of this new technology tomorrow.

Works Cited

Aisch, Gregor, Enrico Bertini, and Moritz Stefaner. Data Stories: Episode 7 | Color. Podcast, 2012. https://datastori.es/episode-7-color-feat-gregor-aisch/.

Bateman, Scott, Regan L. Mandryk, Carl Gutwin, Aaron Genest, David McDine, and Christopher Brooks. "Useful Junk?: The Effects of Visual Embellishment on Comprehension and Memorability of Charts." In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 2573–2582. CHI '10. New York, NY, USA: ACM, 2010. https://doi.org/10.1145/1753326.1753716.

Belcher, Daniel, Mark Billinghurst, S.E. Hayes, and Randy Stiles. "Using Augmented Reality for Visualizing Complex Graphs in Three Dimensions." In *The Second IEEE and ACM International Symposium on Mixed and Augmented Reality, 2003. Proceedings.*, 84–93, 2003. https://doi.org/10.1109/ISMAR.2003.1240691.

Berinato, Scott. Good Charts: The HBR Guide to Making Smarter, More Persuasive Data Visualizations. Boston, Massachusetts: Harvard Business Review Press, 2016.

Bleisch, Susanne. "Toward Appropriate Representations of Quantitative Data in Virtual Environments." *Cartographica* 46, no. 4 (Winter 2011): 252–61. https://doi.org/10.3138/carto.46.4.252.

Bleisch, Susanne, and Jason Dykes. "Quantitative Data Graphics in 3D Desktop-Based Virtual Environments – an Evaluation." *International Journal of Digital Earth* 8, no. 8 (August 3, 2015): 623–39. https://doi.org/10.1080/17538947.2014.927536.

Bleisch, Susanne, Jason Dykes, and Stephan Nebiker. "Evaluating the Effectiveness of Representing Numeric Information Through Abstract Graphics in 3D Desktop Virtual Environments." *The Cartographic Journal* 45, no. 3 (August 2008): 216–26. https://doi.org/10.1179/000870408X311404.

Borgo, R., A. Abdul-Rahman, F. Mohamed, P. W. Grant, I. Reppa, L. Floridi, and Min Chen. "An Empirical Study on Using Visual Embellishments in Visualization." *IEEE Transactions on Visualization and Computer Graphics* 18, no. 12 (December 2012): 2759–68. https://doi.org/10.1109/TVCG.2012.197.

Borkin, M. A., A. A. Vo, Z. Bylinskii, P. Isola, S. Sunkavalli, A. Oliva, and H. Pfister. “What Makes a Visualization Memorable?” *IEEE Transactions on Visualization and Computer Graphics* 19, no. 12 (December 2013): 2306–15. https://doi.org/10.1109/TVCG.2013.234.

Büschel, Wolfgang, Patrick Reipschläger, Ricardo Langner, and Raimund Dachselt. “Investigating the Use of Spatial Interaction for 3D Data Visualization on Mobile Devices.” In *Proceedings of the Interactive Surfaces and Spaces on ZZZ - ISS ’17*, 62–71. Brighton, United Kingdom: ACM Press, 2017. https://doi.org/10.1145/3132272.3134125.

Buschel, Wolfgang, Stefan Vogt, and Raimund Dachselt. “Augmented Reality Graph Visualizations.” *IEEE Computer Graphics & Applications* 39, no. 3 (June 5, 2019): 29–40. https://doi.org/10.1109/MCG.2019.2897927.

Büschel, Wolfgang, Stefan Vogt, and Raimund Dachselt. “Investigating Link Attributes of Graph Visualizations in Mobile Augmented Reality,” n.d., 6.

Butz, Andreas, ed. Smart Graphics: 10th International Symposium, SG 2009, Salamanca, Spain, May 28-30, 2009 ; Proceedings. Lecture Notes in Computer Science 5531. Berlin: Springer, 2009.

Caudell, T. P., and D. W. Mizell. “Augmented Reality: An Application of Heads-up Display Technology to Manual Manufacturing Processes.” In *Proceedings of the Twenty-Fifth Hawaii International Conference on System Sciences*, ii:659–69 vol.2, 1992. https://doi.org/10.1109/HICSS.1992.183317.

Chen, Chaomei, and Yue Yu. “Empirical Studies of Information Visualization: A Meta-Analysis.” *International Journal of Human-Computer Studies* 53, no. 5 (November 2000): 851–66. https://doi.org/10.1006/ijhc.2000.0422.

Chun, Russell. “Redundant Encoding in Data Visualizations: Assessing Perceptual Accuracy and Speed.” *Visual Communication Quarterly* 24, no. 3 (July 2017): 135–48. https://doi.org/10.1080/15551393.2017.1343153.

Cleveland, William S., and Robert McGill. “Graphical Perception: Theory, Experimentation, and Application to the Development of Graphical Methods.” *Journal of the American Statistical Association* 79, no. 387 (1984): 531–54. https://doi.org/10.2307/2288400.

Cohen, I. Bernard. “Florence Nightingale.” *Scientific American* 250, no. 3 (1984): 128–37.

Dübel, S., M. Röhlig, H. Schumann, and M. Trapp. “2D and 3D Presentation of Spatial Data: A Systematic Review.” In *2014 IEEE VIS International Workshop on 3DVis (3DVis)*, 11–18, 2014. https://doi.org/10.1109/3DVis.2014.7160094.

Evergreen, Stephanie, and Chris Metzner. “Design Principles for Data Visualization in Evaluation: Design Principles for Data Visualization.” *New Directions for Evaluation* 2013, no. 140 (December 2013): 5–20. https://doi.org/10.1002/ev.20071.

Experience, World Leaders in Research-Based User. “Why You Only Need to Test with 5 Users.” Nielsen Norman Group. Accessed November 10, 2019. https://www.nngroup.com/articles/why-you-only-need-to-test-with-5-users/.

Freitas, Carla M. D. S., Marcelo S. Pimenta, and Dominique L. Scapin. “User-Centered Evaluation of Information Visualization Techniques: Making the HCI-InfoVis Connection Explicit.” In *Handbook of Human Centric Visualization*, edited by Weidong Huang, 315–36. New York, NY: Springer New York, 2014. https://doi.org/10.1007/978-1-4614-7485-2_12.

Harrison, L., F. Yang, S. Franconeri, and R. Chang. “Ranking Visualizations of Correlation Using Weber’s Law.” *IEEE Transactions on Visualization and Computer Graphics* 20, no. 12 (December 2014): 1943–52. https://doi.org/10.1109/TVCG.2014.2346979.

Heer, Jeffrey, and Michael Bostock. “Crowdsourcing Graphical Perception: Using Mechanical Turk to Assess Visualization Design.” In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, 203–212. CHI ’10. New York, NY, USA: ACM, 2010. https://doi.org/10.1145/1753326.1753357.

Herman, Lukáš, Zdeněk Stachoň, Jan Russnák, Tomáš Řezník, Vojtěch Juřík, and Daniel Vrbík. “Evaluation of User Performance in Interactive and Static 3D Maps.” *ISPRS International Journal of Geo-Information* 7, no. 11 (November 2018): 415. https://doi.org/10.3390/ijgi7110415.

Howard, Alexander Benjamin. “The Art and Science of Data-Driven Journalism,” 2014. https://doi.org/10.7916/D8Q531V1.

Kim, Sung-Hee, Ji Soo Yi, and Niklas Elmqvist. “Oopsy-Daisy: Failure Stories in Quantitative Evaluation Studies for Visualizations.” In *Proceedings of the Fifth Workshop on Beyond Time and Errors Novel*

Evaluation Methods for Visualization - BELIV '14, 142–46. Paris, France: ACM Press, 2014. https://doi.org/10.1145/2669557.2669576.

Lam, H., E. Bertini, P. Isenberg, C. Plaisant, and S. Carpendale. "Empirical Studies in Information Visualization: Seven Scenarios." *IEEE Transactions on Visualization and Computer Graphics* 18, no. 9 (September 2012): 1520–36. https://doi.org/10.1109/TVCG.2011.279.

Lind, Mats, Geoffrey P. Bingham, and Camilla Forsell. "Metric 3D Structure in Visualizations:" *Information Visualization*, July 21, 2016. https://doi.org/10.1057/palgrave.ivs.9500038.

Merritt, Richard K. "Memory Arts to the New Code Paradigm: The Artist as Engineer of Virtual Information Space and Virtual Experience." *Leonardo* 34, no. 5 (December 2001): 403–8. https://doi.org/10.1162/002409401753521511.

Morgan-Wall, Tyler. "Introducing 3D GGPlots with Rayshader," n.d. https://www.tylermw.com/3d-ggplots-with-rayshader/.

Nielsen, Jakob. "First Rule of Usability? Don't Listen to Users." Nielsen Norman Group. Accessed November 10, 2019. https://www.nngroup.com/articles/first-rule-of-usability-dont-listen-to-users/.

———. "Quantitative Studies: How Many Users to Test?" Nielsen Norman Group. Accessed November 10, 2019. https://www.nngroup.com/articles/quantitative-studies-how-many-users/.

———. "Usability 101: Introduction to Usability." Nielsen Norman Group. Accessed November 10, 2019. https://www.nngroup.com/articles/usability-101-introduction-to-usability/.

Nunes, Ana Cecília B. "WSJ R&D: Bringing Data Science and AI Capabilities to the Newsroom." *Global Alliance for Media Innovation* (blog). Accessed November 26, 2019. https://media-innovation.news/media-lab/wsj-rd-bringing-data-science-and-ai-capabilities-to-the-newsroom/.

Pandey, Anshul Vikram, Anjali Manivannan, Oded Nov, Margaret Satterthwaite, and Enrico Bertini. "The Persuasive Power of Data Visualization." *IEEE Transactions on Visualization and Computer Graphics* 20, no. 12 (December 31, 2014): 2211–20. https://doi.org/10.1109/TVCG.2014.2346419.

Pavlik, John V., and Frank Bridges. "The Emergence of Augmented Reality (AR) as a Storytelling Medium in Journalism:" *Journalism & Communication Monographs*, January 13, 2013. https://doi.org/10.1177/1522637912470819.

Peddie, Jon. Augmented Reality: Where We Will All Live. Cham: Springer, 2017.

Porter, Lana Z. “Documenting the World in 3D.” *The New York Times - Research & Development* (blog), November 14, 2019. https://rd.nytimes.com/projects/documenting-the-world-in-3d.

Rangecroft, Margaret. “As Easy as Pie.” *Behaviour & Information Technology* 22, no. 6 (December 11, 2003): 421–26. https://doi.org/10.1080/01449290310001615437.

Rensink, Ronald A., and Gideon Baldridge. “The Perception of Correlation in Scatterplots.” *Computer Graphics Forum* 29, no. 3 (2010): 1203–10. https://doi.org/10.1111/j.1467-8659.2009.01694.x.

Roberts, Graham. “Augmenting Journalism: Immersive Storytelling at The New York Times.” presented at the AWE NYC Meetup, Brooklyn, New York, July 23, 2019. https://www.facebook.com/AugmentedWorldExpo/videos/2328506167403326/.

Roberts, Graham, and Jonathan Corum. “How We Augmented Our Original Reporting of the Moon Landing for Its 50th Anniversary.” *The New York Times*, July 19, 2019, sec. Reader Center. https://www.nytimes.com/2019/07/19/reader-center/apollo-11-moon-landing-augmented-reality.html.

Seipel, Stefan, and Leonor Carvalho. “Solving Combined Geospatial Tasks Using 2D and 3D Bar Charts.” In *2012 16th International Conference on Information Visualisation*, 157–63. Montpellier, France: IEEE, 2012. https://doi.org/10.1109/IV.2012.36.

Shah, Priti, Richard E. Mayer, and Mary Hegarty. “Graphs as Aids to Knowledge Construction: Signaling Techniques for Guiding the Process of Graph Comprehension.” *Journal of Educational Psychology* 91, no. 4 (December 1999): 690–702. https://doi.org/10.1037/0022-0663.91.4.690.

Shepherd, Ifan D. H. “Travails in the Third Dimension: A Critical Evaluation of Three-Dimensional Geographical Visualization.” In *Geographic Visualization*, edited by Martin Dodge, Mary McDerby, and Martin Turner, 199–222. Chichester, UK: John Wiley & Sons, Ltd, 2008. https://doi.org/10.1002/9780470987643.ch10.

Stasko, John. “Value-Driven Evaluation of Visualizations.” In Proceedings of the Fifth Workshop on Beyond Time and Errors Novel Evaluation Methods for Visualization - BELIV ’14, 46–53. Paris, France: ACM Press, 2014. https://doi.org/10.1145/2669557.2669579.

Tankard, Jr., James W. “Effects of Chartoons & Three-Dimensional Graphs on Interest & Information Gain.” *Newspaper Research Journal* 10, no. 3 (Spring 1989): 91–103. https://doi.org/10.1177/073953298901000309.

Tufte, Edward R. *Beautiful Evidence*. Cheshire, Conn: Graphics Press, 2006.

———. *Envisioning Information*. Fourteenth printing. Cheshire, Connecticut: Graphics Press, 2013.

———. *The Visual Display of Quantitative Information*. 2nd ed. Cheshire, Conn: Graphics Press, 2001.

———. Visual Explanations: Images and Quantities, Evidence and Narrative. Cheshire, Conn: Graphics Press, 1997.

Usher, Nikki. Interactive Journalism: Hackers, Data, and Code. Urbana: University of Illinois Press, 2016.

Vovk, Alla, Fridolin Wild, Will Guest, and Timo Kuula. "Simulator Sickness in Augmented Reality Training Using the Microsoft HoloLens." 2018 CHI Conference on Human Factors in Computing Systems, Montreal, Canada, 2018. https://doi.org/10.1145/3173574.3173783.

Wall, E., M. Agnihotri, L. Matzen, K. Divis, M. Haass, A. Endert, and J. Stasko. "A Heuristic Approach to Value-Driven Evaluation of Visualizations." *IEEE Transactions on Visualization and Computer Graphics* 25, no. 1 (January 2019): 491–500. https://doi.org/10.1109/TVCG.2018.2865146.

Yang, Yalong, Bernhard Jenny, Tim Dwyer, Kim Marriott, Haohui Chen, and Maxime Cordeil. "Maps and Globes in Virtual Reality." *Computer Graphics Forum* 37, no. 3 (June 2018): 427–38. https://doi.org/10.1111/cgf.13431.

Zhu, Ying. "Measuring Effective Data Visualization." In *Advances in Visual Computing*, edited by George Bebis, Richard Boyle, Bahram Parvin, Darko Koracin, Nikos Paragios, Syeda-Mahmood Tanveer, Tao Ju, et al., 652–61. Lecture Notes in Computer Science. Springer Berlin Heidelberg, 2007.

Ziemkiewicz, C., and R. Kosara. "The Shaping of Information by Visual Metaphors." *IEEE Transactions on Visualization and Computer Graphics* 14, no. 6 (November 2008): 1269–76. https://doi.org/10.1109/TVCG.2008.171.

Appendix A: Study Consent Form

University of North Carolina at Chapel Hill
IRB Study #: 20-034
Principal Investigator: Peter Andringa

The purpose of this research study is to understand user interactions with Augmented Reality and evaluate potential uses for AR in news data visualizations. You are being asked to take part in this research study because you volunteered through the Hussman School of Journalism and Media's student research participation pool.

Being in a research study is completely voluntary. You can choose not to be in this research study. You can also say yes now and change your mind later.

If you agree to take part in this research, you will be asked to interact with a mobile application using a device we will provide. Using the application, you will explore three different data visualizations through Augmented Reality, and then answer a few survey and interview questions about your experience. Your participation in this study will take about 20 minutes. We expect that about 20 people will take part in this research study.

You can choose not to answer any survey or interview question you do not wish to answer. You can also choose to leave the study at any time. You must be at least 18 years old to participate. If you are younger than 18 years old, please stop now.

The possible risks to you in taking part in this research are:

- A small number of users report feeling mild motion sickness when using immersive applications like Augmented Reality. Because the AR in this project is relatively simple, motion sickness is especially unlikely. Even so, if you start to feel sick at any time, you may stop.

The possible benefits to you for taking part in this research are:

- Experiencing and gaining literacy exploring augmented reality, a new technology that may become more widespread in the future.
- Learning about current events through the data visualizations presented in the study.

To protect your identity as a research subject, the researcher(s) will not share your information with anyone. Those results and recordings will be kept separate from your name, to ensure anonymity. In any publication about this research, your name or other private information will not be used.

If you have any questions about this research, please contact the Investigator named at the top of this form by emailing peter.andringa@unc.edu. If you have questions or concerns about your rights as a research subject, you may contact the UNC Institutional Review Board at 919-966-3113 or by email to IRB_subjects@unc.edu.

Participant name: ________________________ Signature: _______________________

This project was determined to be exempt from federal human subjects research regulations.

Appendix B: Background Survey

Please answer honestly, but don't stress about your responses — just write down or circle whatever comes to mind first. All your responses will be anonymous.

1. **Have you ever used virtual reality? (circle one)**

   Yes　　　　No

2. **Have you ever used augmented reality? (circle one)**

   Yes　　　　No

3. **Define "augmented reality" in a few words or a short sentence:**

4. **Which is the most popular type of augmented reality tool? (Feel free to guess.)**

   A. Large Goggles
   B. Small Glasses
   C. Helmet / Headset
   D. Personal Computer
   E. Smartphone

5. **Circle all of these companies or brands that you've heard of before:**

   HoloLens　　Oculus Rift　　Magic Leap　　Google Glass

6. **Circle which of these apps you have on your phone, if any.**

   Twitter　　New York Times　　Snapchat　　Washington Post

   Instagram　　Facebook　　Wall Street Journal　　TikTok

7. **How often do you read the news?**

   A. Multiple times daily
   B. Daily
   C. Multiple times weekly
   D. Weekly
   E. Less than weekly

8. **How often do you read political news?**

   1 (Rarely)　　2　　3　　4　　5 (All the time)

9. How much do you work with, or read about, data or statistics?

1 (Rarely) 2 3 4 5 (All the time)

10. Are you interested in politics?

1 (Not very) 2 3 4 5
(Extremely)

11. Are you interested in technology?

1 (Not very) 2 3 4 5
(Extremely)

12. What type of smartphone do you use most often?

A. iPhone

B. Android

C. Other

D. I don’t use a smartphone

Return this sheet to the researcher once you’re done.

Appendix C: Screenshots of the AR-Enabled Article

Cover

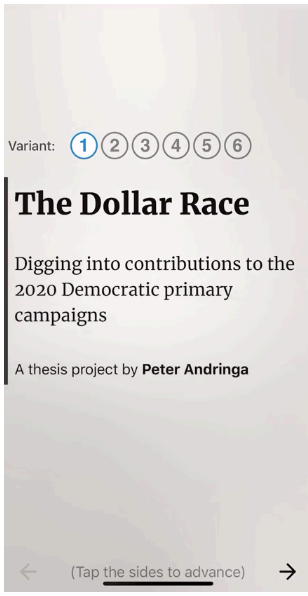

Introduction

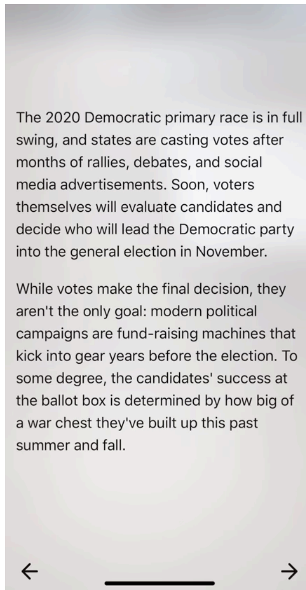

Introduction (cont.)

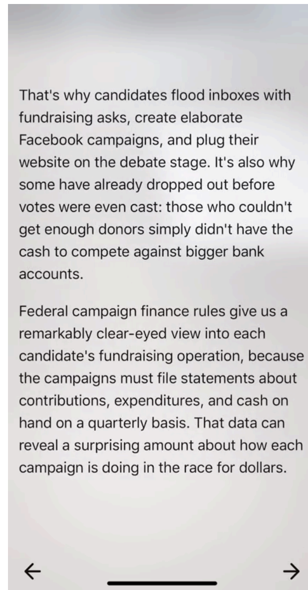

Bar Chart intro

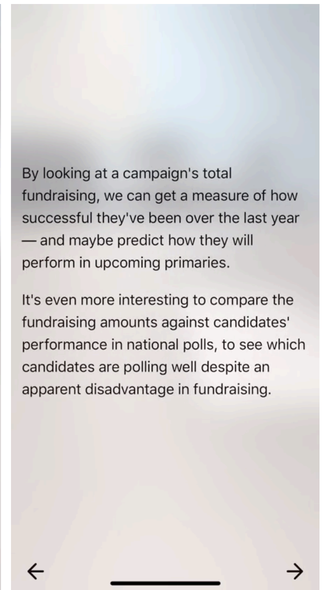

Bar Chart AR

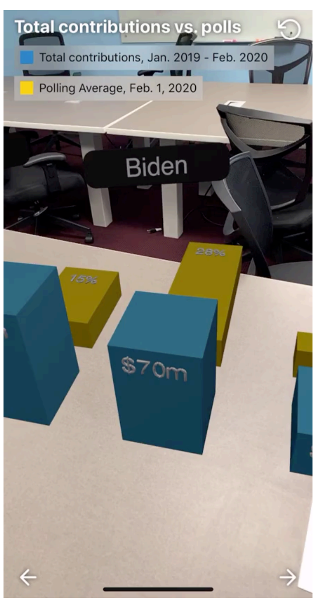

Map intro

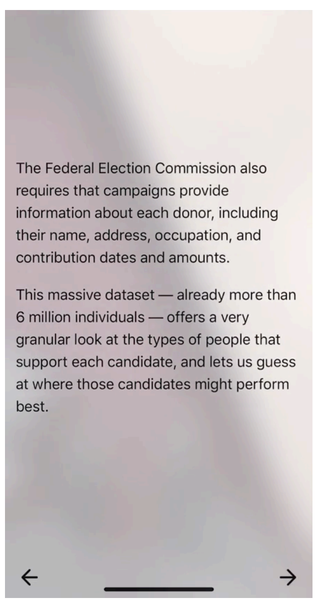

Map AR (1/3)

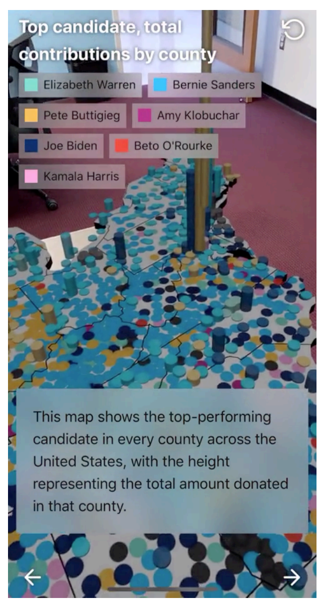

Map AR (2/3)

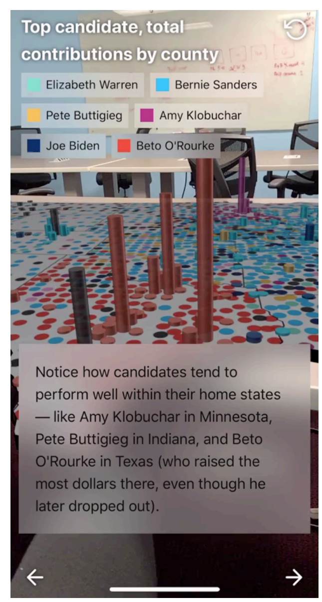

Map AR (3/3)

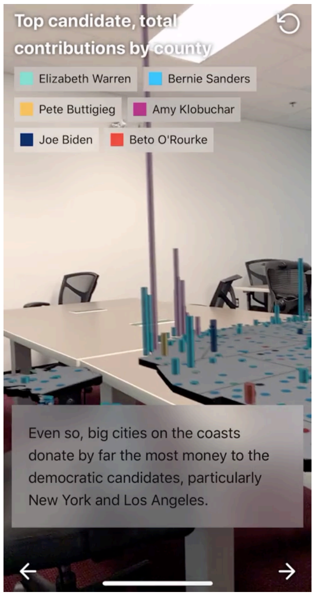

Spheres intro

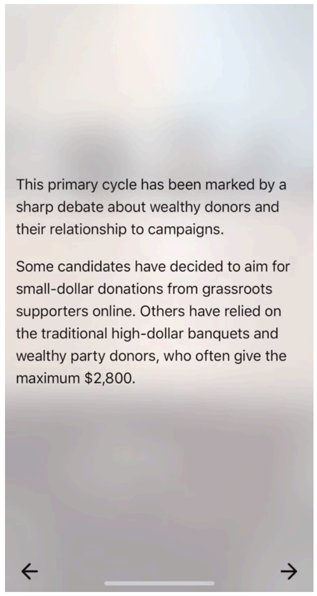

Spheres intro (cont.)

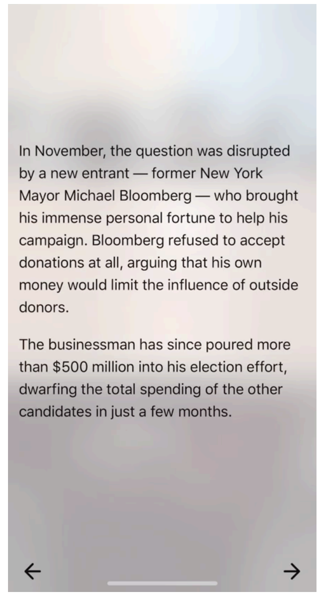

Spheres AR

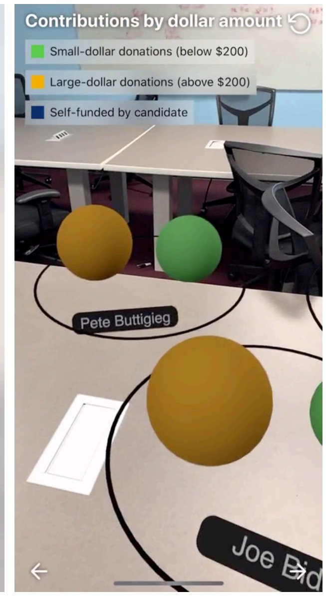

Conclusion

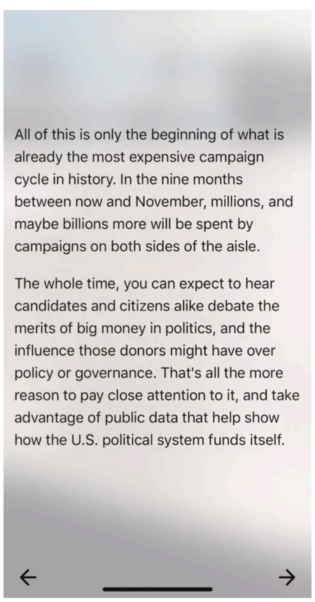

Credits

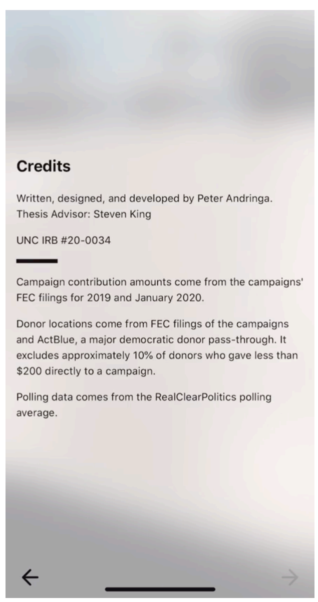

Note: The AR sections (and their associated introductory text) were shuffled randomly for each user, using one of the "Variant" selectors on the cover page.

Appendix D: Post-Experience Survey

Please answer honestly, but don't stress about your responses — just write down or circle whatever comes to mind first. All your responses will be anonymous.

13. What was one thing you learned from the article?

14. Of the three graphics, which visualization felt most interesting? Why?

Questions about the bar chart of candidate fundraising and polling:

15. Based on the graphic, which candidate had raised the most money overall?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

16. Based on the graphic, which candidate has raised the least?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

17. For the following questions, circle the candidate who raised <u>more money</u>:

A.	Elizabeth Warren	*or*	Joe Biden
B.	Pete Buttigieg	*or*	Bernie Sanders
C.	Joe Biden	*or*	Amy Klobuchar
D.	Bernie Sanders	*or*	Elizabeth Warren
E.	Michael Bloomberg	*or*	Tom Steyer

18. Based on the graphic, which candidate was polling the highest?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

19. Overall, how clear and easy to understand was this graphic?

1 (very unclear) 2 3 4 5 (very clear)

Questions about the map of donations:

20. Which candidate raised the most near Raleigh / Chapel Hill?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

21. Which candidate appeared to have the broadest nationwide support?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

22. Which of these candidates had the most concentrated support from just one location?

A. Elizabeth Warren
B. Amy Klobuchar
C. Bernie Sanders
D. Tom Steyer

23. Which of these candidates was <u>not</u> shown on the map?

A. Kamala Harris
B. Andrew Yang
C. Joe Biden
D. Beto O'Rourke

9. Write one fact that stuck out to you when you looked at the map:

24. Overall, how clear and easy to understand was this map?

1 (very unclear) 2 3 4 5 (very clear)

Questions about the graphic showing spheres of donation types:

25. Which candidate has a larger proportion of small donors?

A.	Elizabeth Warren	*or*	Joe Biden
B.	Pete Buttigieg	*or*	Bernie Sanders
C.	Joe Biden	*or*	Amy Klobuchar
D.	Tom Steyer	*or*	Elizabeth Warren
E.	Michael Bloomberg	*or*	Tom Steyer

26. Which candidate had the most different results from the others?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

27. Which candidate had raised the most money?

A. Joe Biden
B. Amy Klobuchar
C. Michael Bloomberg
D. Pete Buttigieg
E. Elizabeth Warren
F. Bernie Sanders
G. Andrew Yang
H. Tom Steyer

28. What was the cut-off between a "small" donor and a "large" one?

A. $100
B. $200
C. $300
D. $500
E. $1000
F. $2800

29. Overall, how clear and easy to understand was this graphic?

1 (very unclear) 2 3 4 5 (very clear)

Return this packet to the researcher once you're done.

Appendix E: In-Person Interview Questions

These questions were asked by the interviewer after the subject completed the written post-experience survey (in Appendix B). They were designed to fill in any gaps in the observation notes and elicit more explanation about how the user felt.

1. What part of the experience felt the most natural for you to use?
2. What part of the experience felt the most unnatural?
3. Did you ever feel confused? During which section?
4. Did anything surprise you? Why?
5. Do you feel like this was more or less useful than traditional 2D graphics? Why?
6. Have you used AR before? Do you think you will use it more after trying it here?
7. Do you think you would read more news stories in this format, if they were available? Why or why not?

Appendix F: A (partial) list of AR stories from major U.S. publications

The New York Times (https://www.nytimes.com/spotlight/augmented-reality**)**

“Apollo 11: As They Shot It.” July 18, 2019. (https://www.nytimes.com/interactive/2019/07/18/science/apollo-11-moon-landing-photos-ul.html)

“Why Notre-Dame was a Tinderbox.” April 17, 2019. (https://www.nytimes.com/interactive/2019/04/17/world/europe/notre-dame-cathedral-fire-spread.html)

“A Closer Look at the Polar Vortex’s Dangerously Cold Winds.” January 30, 2019. (https://www.nytimes.com/interactive/2019/01/30/science/polar-vortex-extreme-cold.html)

“It’s Intermission for the Large Hadron Collider.” December 21, 2018. (https://www.nytimes.com/interactive/2018/12/21/science/cern-large-hadron-collider-ar-ul.html)

“Lakeith Stanfield’s Balancing Act.” December 5, 2018. (https://www.nytimes.com/interactive/2018/12/05/magazine/lakeith-stanfield-great-performers-ar-ul.html)

“Take a Tour of Lady Liberty’s Torch (Right This Second).” November 13, 2018. (https://www.nytimes.com/interactive/2018/11/13/nyregion/statue-of-liberty-torch-ar-ul.html)

“Ashley Graham, Unfiltered.” September 5, 2018. (https://www.nytimes.com/interactive/2018/09/04/style/ashley-graham-body-positive-movement-ar-ul.html)

"Step Inside the Thai Cave in Augmented Reality.” July 21, 2018. (https://www.nytimes.com/interactive/2018/07/21/world/asia/thai-cave-rescue-ar-ul.html)

“How We Created a Virtual Crime Scene to Investigate Syria’s Chemical Attack.” June 24, 2018. (https://www.nytimes.com/interactive/2018/06/24/world/middleeast/douma-syria-chemical-attack-augmented-reality-ar-ul.html)

"Augmented Reality: David Bowie in Three Dimensions." March 20, 2018. (https://www.nytimes.com/interactive/2018/03/20/arts/design/bowie-costumes-ar-3d-ul.html)

"Augmented Reality: Four of the Best Olympians, as You've Never Seen Them." February 5, 2018. (https://www.nytimes.com/interactive/2018/02/05/sports/olympics/ar-augmented-reality-olympic-athletes-ul.html)

The Washington Post

"How to dress for space." July 17, 2019. (https://www.washingtonpost.com/graphics/2019/business/immersive-space-suits-history-fashion-and-function/)

"What remains of Bears Ears." April 2, 2019. (https://www.washingtonpost.com/graphics/2019/national/bears-ears/)

"Use augmented reality stencils this year to carve your pumpkin." October 26, 2018. (https://www.washingtonpost.com/lifestyle/use-augmented-reality-stencils-this-year-to-carve-your-pumpkin/2018/10/26/d1900e16-d866-11e8-aeb7-ddcad4a0a54e_story.html)

"Practice makes perfect: Carve this virtual turkey." November 16, 2018. (https://www.washingtonpost.com/graphics/2018/food/how-to-carve-a-turkey-augmented-reality/)

"A mystery dinosaur in the basement." September 18, 2018. (https://www.washingtonpost.com/graphics/2018/national/smithsonian-dinosaur-augmented-reality/)

"World-class heritage – here in the U.S." May 3, 2018. (https://www.washingtonpost.com/graphics/2018/lifestyle/travel/us-unesco-guide/)

"The speediest Winter Olympic sports." February 2, 2018. (https://www.washingtonpost.com/graphics/2018/sports/fastest-winter-sports/)

“The Smithsonian saved a statue Lincoln praised. Now you can judge it for yourself.” October 27, 2017. (https://www.washingtonpost.com/graphics/2017/national/photogrammetry-civil-war-statuette/)

“The African American Museum: Still the hottest ticket in town.” September 22, 2017. (https://www.washingtonpost.com/entertainment/museums/the-african-american-museum-a-year-later-still-the-hottest-ticket-in-town/2017/09/22/15aa027e-8cf4-11e7-84c0-02cc069f2c37_story.html

The Wall Street Journal

"Experience augmented reality in the WSJ app for iOS.” September 21, 2017. (https://www.dowjones.com/press-room/experience-augmented-reality-wsj-app-ios/)

Time Magazine

“Landing on the Moon: Apollo 11.” July 18, 2019. (https://time.com/longform/apollo-11-moon-landing-immersive-experience/)

“Inside the Amazon: The Dying Forest.” September 12, 2019. (https://time.com/longform/inside-amazon-rain-forest-vr-app/)

Quartz

“The 2050 Project.” March 2019. (https://qz.com/se/the-2050-project/)

The New Yorker

"Introducing New Yorker Cartoons in Augmented Reality.” November 25, 2019. (https://www.newyorker.com/cartoons/cartoon-desk/introducing-new-yorker-cartoons-in-augmented-reality)